A Duality Based Approach for Network Revenue Management in Airline Alliances

Huseyin Topaloglu
School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853, USA
topaloglu@orie.cornell.edu

December 27, 2011

Abstract

In this paper, we develop a revenue management model for airline alliances. In particular, our model allows the members of the alliance to decide how to share the revenues from itineraries spanning flight legs operated by multiple airlines and how to coordinate their capacity allocation decisions. We start with a deterministic linear programming approximation to the revenue management problem that takes place over the whole alliance network. This linear program models the decisions made by a central planner. In the linear program, we relax the constraints that link the decisions for different airlines by associating dual multipliers with them, in which case, the linear program decomposes by the airlines into smaller linear programs. We extract a booking control policy from the smaller linear programs that allows each airline to make its booking control decisions autonomously. Furthermore, the dual multipliers indicate how to share the revenues from itineraries that span flight legs operated by multiple airlines. We prove that the revenue sharing scheme provided by our approach has a number of intuitively appealing features. Computational experiments indicate that our model and our revenue sharing scheme coordinate the decisions of the alliance members quite well.

1 Introduction

Building alliances represents one of the most vital tactics for increasing market share and ensuring profitability for airlines. Through codeshare agreements, members of an alliance gain access to the inventory of seats available on the flight legs operated by the other airlines in the alliance, allowing them to serve origin destination markets that would be impossible to serve by using only their own networks. Beside broadening the origin destination markets served by an airline, codeshare agreements ensure that the flight legs operated by an airline pool the demand from multiple members of the alliance, resulting in economies of scale. While such benefits of alliances are undeniable, building and operating an alliance bring a number of challenges from revenue management perspective. When an airline markets an itinerary that uses flight legs operated by other airlines in the alliance, the alliance members need to agree on how to share the revenue obtained from the itinerary. Similarly, since an accepted itinerary request may consume capacities on the flight legs operated by multiple airlines, the revenue management decisions of the airlines in the alliance interact with each other, creating a need for the alliance members to coordinate their booking control policies.

In this paper, we develop an alliance revenue management model that allows the airlines to decide how share the revenues from itineraries that use flight legs operated by multiple airlines and how to coordinate their capacity allocation decisions. We consider an alliance where each member operates a set of flight legs and markets a set of itineraries. When an airline accepts a request for an itinerary that it markets, it generates a revenue and this revenue is shared among the airline marketing the itinerary and the airlines operating the flight legs used by the itinerary. The goal of each airline is to maximize its total expected revenue. On the other hand, from a system design perspective, the goal is to ensure that each airline is content to be a part of the alliance.

The starting point for our model is a deterministic linear programming approximation to the revenue management problem that takes place over the airline network operated by the whole alliance. This deterministic linear program models the booking control decisions made by a central planner overseeing the alliance. In the deterministic linear program, there are decision variables representing the itinerary acceptance decisions for the itineraries marketed by different airlines. We decompose the deterministic linear program by associating dual multipliers with the constraints that link the decisions for different airlines. In this case, we obtain a number of smaller linear programs, each of which corresponds to a different airline in the alliance. Each one of the smaller linear programs corresponding to a particular airline involves only the itineraries marketed by this airline and the flight legs operated by this airline. We show how to extract a booking control policy from the smaller linear programs that allows each airline to make its capacity allocation decisions autonomously. Furthermore, the dual multipliers that we use to decompose the original deterministic linear program provide fare allocations, indicating how the revenue from an itinerary should be shared among the alliance members.

Our model makes several practical and technical contributions. On the practical side, the booking control policy we propose allows each airline to make its capacity control decisions autonomously, while considering its impact on the whole alliance. The booking control policy used by our model is a bid price policy. The idea behind a bid price policy is to associate bid prices with the flight legs, capturing the opportunity cost of a seat on different flight legs. In this case, an itinerary request is accepted only if the revenue associated with the itinerary exceeds the total opportunity cost of the seats consumed by the itinerary request. In our model, each alliance member can compute its bid prices by using the remaining capacities only on the flight legs that it operates and the expected demands only for the itineraries that it markets. This feature is practically useful since airlines recompute bid prices over the selling horizon in response to changing remaining capacities and changing expected demands and it would be difficult to frequently access the remaining capacity and expected demand information from other airlines. In our model, the only time the airlines need to share information is when they reach an agreement on a revenue sharing scheme at the beginning of the selling horizon. Overall, our computational experiments indicate that although the alliance members act autonomously, the alliance suffers 2.76% revenue loss on average when compared with a central planner that can oversee the whole network.

On the technical side, the revenue sharing scheme stipulated by our model has a number of interesting features. To begin with, we show that if an airline does not market a particular itinerary, then its revenue share of the itinerary should simply be equal to the total opportunity cost of the seats consumed by the itinerary on the flight legs operated by this airline (Proposition 1). In the revenue management literature, the revenue shares are anecdotally linked to the opportunity costs of the seats and there are some small dynamic programming models that justify this intuition, but our model provides a justification for this intuition for large airline networks. Furthermore, our revenue sharing scheme satisfies the natural expectation that if an itinerary is not marketed by a particular airline and does not use any of the flight legs operated by the airline, then this airline should not generate any revenue from the sale of the itinerary (Proposition 2). Another interesting aspect of our revenue sharing scheme appears when we compare the revenue shares of two itineraries. Consider two itineraries that are not marketed by a particular airline, but they consume the same capacities on the flight legs operated by this airline. In this case, the revenue shares of these itineraries for the particular airline are identical (Proposition 3). Therefore, although the original fares associated with the two itineraries can be quite different, the revenue benefits of the nonmarketing airline from these itineraries are not different at all. Finally, it turns out that if the airlines compute their bid prices once at the beginning and do not recompute them during the course of the selling horizon, then the autonomous capacity allocation decisions made by each airline would be identical to the capacity allocation decisions made by a central planner (Proposition 4). This result provides theoretical support for the autonomous booking control policy that we develop for each airline and indicates that the airlines do not have incentive to deviate from the decisions made by a central planner, at least until they recompute their bid prices.

The rest of the paper is organized as follows. In Section 2, we review the related literature. In Section 3, we describe the main components of our alliance model and our notation. In Section 4, we describe the deterministic linear programming approximation that models the capacity allocation decisions made by a central planner. In Section 5, we show how to decompose the deterministic linear program by the airlines. The duality idea used in the decomposition process provides a revenue sharing scheme among the airlines. In Section 6, we show how our model allows each airline to make its booking control decisions autonomously. In Section 7, we study the properties of the revenue sharing mechanism from our model. In Section 8, we present computational experiments. In Section 9, we conclude.

2 Literature Review

Revenue management models for airline alliances form an active area of research. Boyd (1998) describes a mathematical programming model to allocate the capacities on the flight legs among the different members of an alliance. His model uses the same deterministic linear program we build on, but he does not describe how to set up a revenue sharing scheme or how to make alliance revenue management decisions. Without a mathematical model, Vinod (2005) intuitively hints at the expectation that the revenue sharing scheme of an alliance should be related to bid prices. Houghtalen et al. (2007) build coordination mechanism in an airline cargo alliance and use a capacity allocation approach similar to the one in Boyd (1998). Wright et al. (2010) develop a Markovian game model for an alliance with two airlines. They study revenue sharing schemes that ensure that the total expected revenue of the alliance matches the one obtained by a central planner, but it is computationally difficult to implement their model on large airline networks. Hu et al. (2010) build a game theoretic model of an alliance, where alliance members use partitioned booking limit policies instead of bid price policies. They setup a revenue sharing scheme to coordinate the decisions of the airlines. Wright (2011) studies a number of revenue sharing schemes to coordinate the decisions of alliance members and his work can be interpreted as an extension of Wright et al. (2010) to deal with realistic airline networks.

Our alliance model uses a deterministic linear programming approximation to the network revenue management problem as a starting point. This linear program dates back to the work of Simpson (1989) and Williamson (1992) and it is formulated under the assumption that all itinerary requests take on their expected values. Airline industry widely uses this deterministic linear program to construct booking control policies. Since our alliance model rests on such a widely used approximation, we hope that it will quickly be appealing to practitioners. There are several variants of the deterministic linear programming approximation. Talluri and van Ryzin (1999) propose a randomized variant that uses samples of the itinerary requests, rather than expectations. Their goal is to capture the probabilistic nature of the requests more accurately. Adelman (2007) proposes a dynamic variant of the deterministic linear program that more accurately captures the dynamic nature of the arrivals of the itinerary requests over time. Computational experiments indicate that these refinements significantly improve the quality of the booking control policies obtained from the original deterministic linear program. Although we do not pursue these extensions here, it turns out that one can follow the same line of reasoning in this paper to decompose the randomized and dynamic variants of the deterministic linear program by the airlines and extract booking control policies from these variants in an alliance setting.

Decomposition is a central theme in this paper. Decomposition ideas appear in the optimization literature dating back to Dantzig and Wolfe (1960). In revenue management, such ideas are generally motivated by the necessity to decompose large network revenue management problems by the flight legs so that one can obtain high quality policies in a tractable fashion. Liu and van Ryzin (2008), Topaloglu (2009), Erdelyi and Topaloglu (2011) and Zhang (2011) are representative decomposition approaches in network revenue management. In contrast, we decompose the network revenue management problem by the airlines rather than by the flight legs and our interest in decomposition originates from the fact that the decomposed problem more accurately reflects the decision environment of the alliance.

3 Model Primitives

We consider an alliance that includes a number of airlines. Each airline operates a set of flight legs and markets a set of itineraries. An itinerary marketed by a particular airline may use capacities on flight legs operated by multiple airlines. Over time, each airline receives requests for the itineraries that it markets and it needs to decide whether to accept or reject each itinerary request. An accepted itinerary request uses capacities on one or more flight legs and generates a revenue, which is shared among the airlines operating the flight legs used by the itinerary. Therefore, an airline may generate revenue not only from the itineraries that it markets, but also from the itineraries marketed by other airlines as long as these itineraries use some of the flight legs operated by the airline. From an operational perspective, the objective of each airline is to maximize its total expected revenue. On the other hand, from a system design perspective, the objective is to ensure that each airline is content to be a part of the alliance and their incentives are compatible with each other.

We use $\mathcal{K}$ to denote the set of airlines in the alliance. The set of flight legs operated by the alliance is $\mathcal{L}$ and the set of itineraries marketed by the alliance is $\mathcal{J}$. We use $\mathcal{L}^k$ to denote the set of flight legs operated by airline k and $\mathcal{J}^k$ to denote the set of itineraries marketed by airline k. We assume that $\cup_{k \in \mathcal{K}} \mathcal{L}^k = \mathcal{L}$ and $\mathcal{L}^k \cap \mathcal{L}^\ell = \varnothing$ for $k \neq \ell$ so that the sets of flight legs $\{\mathcal{L}^k : k \in \mathcal{K}\}$ operated by the airlines partition the set of flight legs $\mathcal{L}$. Similarly, we assume that $\cup_{k \in \mathcal{K}} \mathcal{J}^k = \mathcal{J}$ and $\mathcal{J}^k \cap \mathcal{J}^\ell = \varnothing$ for $k \neq \ell$. We let $a_{ij} = 1$ if itinerary j uses flight leg i. Otherwise, we let $a_{ij} = 0$. By using the notation defined so far, we can easily deduce which itineraries use flight legs operated by multiple airlines and which itineraries use flight legs operated by a single airline. In particular, the set of flight legs used by itinerary j is given by $\{i \in \mathcal{L} : a_{ij} = 1\}$ so that itinerary j uses flight legs operated by a single airline if and only if all of the flight legs in this set are operated by a single airline. We use c_i to denote the total available capacity on flight leg i.

The fare associated with itinerary j is f_j. We use p_j^k to denote the portion of the fare associated with itinerary j that is allocated to airline k. In other words, if the airline marketing itinerary j accepts a request for this itinerary, then airline k generates a revenue of p_j^k. The fare allocations satisfy $\sum_{k \in \mathcal{K}} p_j^k = f_j$ so that the total fare allocation of itinerary j to all of the airlines is equal to the fare associated with itinerary j. It is also reasonable to have $p_j^k = 0$ whenever itinerary j is not marketed by airline k and does not use any of the flight legs operated by airline k. Therefore, the airlines that do not market a particular itinerary and do not operate any of the flight legs used by this itinerary do not generate any revenue from the sale of the itinerary. We assume that the airline that markets itinerary j decides whether to accept or reject a request for itinerary j. As long as there is enough capacity on all of the flight legs used by itinerary j, the other airlines comply with the marketing airline, implying that they are content with the fare allocations set at the beginning of the booking horizon. This setup is slightly different than the free sale environment in Boyd (1998), where any of the airlines marketing itinerary j or operating the flight legs used by itinerary j can veto a sale for this itinerary.

Our development does not depend heavily on the arrival process for the itinerary requests, but we specify one for concreteness. The requests for itinerary j arrive over the booking horizon $[0, \tau]$ according to a Poisson process with rate function $\{\Lambda_j(t) : t \in [0, \tau]\}$. In this case, the total expected demand for itinerary j is $\lambda_j = \int_0^\tau \Lambda_j(t)\,dt$. In the next section, we begin by giving a relatively simple but realistic model of how a central planner overseeing the whole alliance can estimate the total expected revenue of the alliance and make booking control decisions. It turns out that we can build on this model to come up with fare allocations $\{p_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$, which ultimately allow each airline to solve a separate problem to make its own booking control decisions autonomously. The appealing aspect of the problem solved by each airline k is that this problem requires having access to only the capacities on the flight legs operated by airline k and the demands for the itineraries marketed by airline k.

4 Total Expected Revenue Estimated by Central Planner

If there is a central planner that oversees the whole alliance, then we can try to solve a dynamic program to find the optimal booking control policy for the alliance. However, the state variable in this dynamic program ends up being a high dimensional vector since it has to keep track of the remaining capacity on each flight leg. As a result, solving a dynamic program to compute the optimal booking control policy is usually not possible for realistic airline networks. In practice, a common approach to get around this difficulty is to use a deterministic approximation that is formulated under the assumption that all itinerary requests take on their expected values. In particular, letting x_j be the total number of requests for itinerary j that a central planner would want to accept, we can estimate the total expected revenue of the alliance by solving the linear program

$$\hat{Z} = \max \quad \sum_{j \in \mathcal{J}} f_j\, x_j \tag{1}$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} a_{ij}\, x_j \leq c_i \qquad \forall\, i \in \mathcal{L} \tag{2}$$

$$x_j \leq \lambda_j \qquad \forall\, j \in \mathcal{J} \tag{3}$$

$$x_j \geq 0 \qquad \forall\, j \in \mathcal{J}. \tag{4}$$

The objective function of the problem above accounts for the total revenue. Constraints (2) ensure that the capacity allocation decisions of the central planner do not violate the capacities on the flight legs, whereas constraints (3) ensure that the numbers of requests served for the itineraries do not exceed the total expected demands. Problem (1)-(4) is referred to as the deterministic linear programming approximation to the network revenue management problem and it dates back to the work of Simpson (1989) and Williamson (1992). Letting $\hat{x} = \{\hat{x}_j : j \in \mathcal{J}\}$ be an optimal solution to problem (1)-(4), we can view $\hat{x}$ as the ideal capacity allocation decisions from the perspective of a central planner, yielding an estimated total expected revenue of $\hat{Z}$.

There are two important uses of the deterministic approximation in problem (1)-(4). First, it is possible to show that the optimal objective value of problem (1)-(4) is an upper bound on the total expected revenue collected by any nonanticipatory policy that can be used by a central planner. Such an upper bound on the optimal total expected revenue serves as a useful benchmark when testing the performance of suboptimal control policies. Second, we can utilize problem (1)-(4) to derive a booking control policy that can be used by a central planner. In particular, letting $\{\hat{\mu}_i : i \in \mathcal{L}\}$ be the optimal values of the dual variables associated with constraints (2), we can use $\hat{\mu}_i$ to estimate the value of a unit of capacity on flight leg i. In network revenue management, $\hat{\mu}_i$ is called the bid price of flight leg i. The booking control policy derived from problem (1)-(4) accepts a request for itinerary j if the revenue from itinerary j exceeds the total value of the capacity consumed by this itinerary and there is enough capacity to serve the itinerary request. Thus, we serve a request for itinerary j if

$$f_j \geq \sum_{i \in \mathcal{L}} a_{ij} \, \hat{\mu}_i \tag{5}$$

and there is enough capacity to serve a request for itinerary j. The decision rule in (5) is quite popular in practice. Talluri and van Ryzin (1998) provide theoretical support for this decision rule by showing that it is asymptotically optimal in a regime where the capacities on the flight legs and the expected demand for the itineraries grow linearly with the same rate.

Implementing the decision rule above requires solving problem (1)-(4), which uses the capacities $\{c_i : i \in \mathcal{L}\}$ on all of the flight legs and the expected demands $\{\lambda_j : j \in \mathcal{J}\}$ for all of the itineraries in the alliance. A central planner having access to all of this information generally does not exist in alliances and each airline makes its own decisions somewhat autonomously. In the next section, we decompose problem (1)-(4) by the airlines to come up with fare allocations $\{p_j^k : k \in \mathcal{K}, \ j \in \mathcal{J}\}$. Once we have the fare allocations, we construct a booking control policy for each airline in the alliance by solving a separate problem that uses only the information available to a particular airline.

5 Computing Fare Allocations

In this section, we decompose problem (1)-(4) by the airlines and the decomposed version of problem (1)-(4) eventually becomes useful for coming up with fare allocations. We begin by augmenting the set of airlines by a fictitious airline ψ so that the set of airlines becomes $\mathcal{K} \cup \{\psi\}$. We assume that airline ψ does not operate any flight legs or market any itineraries so that the sets of flight legs and itineraries remain unchanged. In this case, using the decision variables $\{x_j^k : k \in \mathcal{K} \cup \{\psi\}, \ j \in \mathcal{J}\}$ instead of $\{x_j : j \in \mathcal{J}\}$, we can equivalently write problem (1)-(4) as

$$\max \quad \sum_{j \in \mathcal{J}} f_j \, x_j^\psi \tag{6}$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} a_{ij} \, x_j^k \leq c_i \qquad \forall \, k \in \mathcal{K}, \ i \in \mathcal{L}^k \tag{7}$$

$$x_j^k \leq \lambda_j \qquad \forall \, k \in \mathcal{K}, \ j \in \mathcal{J}^k \tag{8}$$

$$x_j^\psi - x_j^k = 0 \qquad \forall \, k \in \mathcal{K}, \ j \in \mathcal{J} \tag{9}$$

$$x_j^k \geq 0 \qquad \forall \, k \in \mathcal{K}, \ j \in \mathcal{J}. \tag{10}$$

To see the equivalence between problems (1)-(4) and (6)-(10), we note that we can use constraints (9) to replace all of the decision variables $\{x_j^k : k \in \mathcal{K}\}$ in problem (6)-(10) with a single decision variable x_j^ψ. In this case, we can drop constraints (9) from problem (6)-(10). Furthermore, since $\cup_{k \in \mathcal{K}} \mathcal{L}^k = \mathcal{L}$ and $\cup_{k \in \mathcal{K}} \mathcal{J}^k = \mathcal{J}$, constraints (7) and (8) enforce the capacity availability for all of the flight legs and the demand availability for all of the itineraries. Therefore, problems (1)-(4) and (6)-(10) are equivalent to each other and recalling the notation in Section 4, the optimal objective value of problem (6)-(10) is still $\hat{Z}$. It is useful to observe that problem (6)-(10) has the block angular structure that is amenable to the decomposition method of Dantzig and Wolfe (1960), where the blocks correspond to the different airlines and constraints (9) play the role of linking the decisions for the different airlines.

To decompose problem (6)-(10) by the airlines, we let $\{\hat{\alpha}_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$ be the optimal values of the dual variables associated with constraints (9) in problem (6)-(10). Dualizing constraints (9) by associating the multipliers $\{\hat{\alpha}_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$ with them, the objective function of problem (6)-(10) becomes $\sum_{j\in\mathcal{J}}[f_j - \sum_{k\in\mathcal{K}} \hat{\alpha}_j^k]\, x_j^\psi + \sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{J}} \hat{\alpha}_j^k\, x_j^k$. By linear programming duality, it follows that maximizing the last expression subject to constraints (7), (8) and (10) is equivalent to solving problem (6)-(10). On the other hand, by the constraint in the dual of problem (6)-(10) associated with the decision variable x_j^ψ, we have $\sum_{k\in\mathcal{K}} \hat{\alpha}_j^k = f_j$, which implies that the term $f_j - \sum_{k\in\mathcal{K}} \hat{\alpha}_j^k$ in the last objective function expression is zero for all $j \in \mathcal{J}$. Therefore, the optimal objective value of problem (6)-(10) is equal to the optimal objective value of the problem

$$\max \quad \sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{J}} \hat{\alpha}_j^k\, x_j^k \tag{11}$$

$$\text{subject to} \quad (7), (8), (10). \tag{12}$$

The crucial observation is that the objective function and all of constraints (7), (8) and (10) in the problem above decomposes by the airlines. This implies that problem (11)-(12) decomposes into $|\mathcal{K}|$ separate problems and the problem for airline k has the form

$$\hat{Z}^k = \max \quad \sum_{j\in\mathcal{J}} \hat{\alpha}_j^k\, x_j^k \tag{13}$$

$$\text{subject to} \quad \sum_{j\in\mathcal{J}} a_{ij}\, x_j^k \le c_i \qquad \forall\, i \in \mathcal{L}^k \tag{14}$$

$$x_j^k \le \lambda_j \qquad \forall\, j \in \mathcal{J}^k \tag{15}$$

$$x_j^k \ge 0 \qquad \forall\, j \in \mathcal{J}. \tag{16}$$

Summing up the discussion so far, noting that the optimal objective value of the problem above is denoted by $\hat{Z}^k$, we have $\sum_{k\in\mathcal{K}} \hat{Z}^k = \hat{Z}$ and letting $\hat{x}^k = \{x_j^k : j \in \mathcal{J}\}$ be the corresponding optimal solution, we can view $\hat{x}^k$ as the ideal capacity allocation decisions from the perspective of airline k.

We interpret $\hat{\alpha}_j^k$ in problem (13)-(16) as the fare allocation of itinerary j to airline k. Noting that $\sum_{k\in\mathcal{K}} \hat{\alpha}_j^k = f_j$ for all $j \in \mathcal{J}$, the total fare allocation of an itinerary to all of the airlines is indeed equal to the fare associated with the itinerary. We note that the structure of problem (13)-(16) is similar to that of problem (1)-(4), but the objective function of problem (13)-(16) estimates the total expected revenue of airline k as given by the fare allocations to this airline, while constraints (14) enforce the capacity availability only on the flight legs operated by airline k and constraints (15) enforce the demand availability only for the itineraries marketed by airline k. Thus, problem (13)-(16) uses only the information that is available to airline k. It is also worthwhile to notice that problem (13)-(16) is a decomposed version of problem (11)-(12) and the latter problem is equivalent to problem (1)-(4). Therefore, noting that $\hat{x}$ is an optimal solution to problem (1)-(4), we obtain an optimal solution to problem (13)-(16) by letting $x_j^k = \hat{x}_j$ for all $j \in \mathcal{J}$. In other words, the fare allocations $\{\hat{\alpha}_j^k : j \in \mathcal{J}\}$

of airline k are such that an optimal solution to problem (13)-(16) is also an optimal solution to problem (1)-(4) solved by a central planner, indicating that the airlines do not have an incentive to deviate from the centrally optimal capacity allocation decisions.

In the next section, we come up with an autonomous booking control policy for each airline.

6 Booking Control by Each Airline

The discussion in the previous section suggests that if we let $\{\hat{\alpha}_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$ be the optimal values of the dual variables associated with constraints (9) in problem (6)-(10), then it is sensible to use $\{\hat{\alpha}_j^k : j \in \mathcal{J}\}$ as the fare allocations to airline k. Obtaining these fare allocations requires solving problem (6)-(10), which, in turn, requires having access to the capacities on all of the flight legs and the expected demands for all of the itineraries marketed by the alliance. Although not always possible, access to this information may be reasonable for the purpose of computing the fare allocations since the fare allocations are computed relatively infrequently. On the other hand, the question we seek to answer in this section is how each airline makes its own booking control decisions in real time with little access to up to date capacity and expected demand information from the other airlines.

For given fare allocations $\{p_j^k : j \in \mathcal{J}\}$ to airline k, which may have been obtained by using an optimal dual solution to problem (6)-(10), problem (13)-(16) motivates making the booking control decisions for this airline by solving the problem

$$\max \quad \sum_{j \in \mathcal{J}} p_j^k \, x_j^k \tag{17}$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} a_{ij} \, x_j^k \le c_i \qquad \forall\, i \in \mathcal{L}^k \tag{18}$$

$$x_j^k \le \lambda_j \qquad \forall\, j \in \mathcal{J}^k \tag{19}$$

$$x_j^k \ge 0 \qquad \forall\, j \in \mathcal{J}. \tag{20}$$

The problem above immediately allows us to construct a booking control policy similar to the one in (5). In particular, if we let $\{\hat{\pi}_i^k : i \in \mathcal{L}^k\}$ be the optimal values of the dual variables associated with constraints (18) in problem (17)-(20), then we can use $\hat{\pi}_i^k$ to estimate the value of a unit of capacity on flight leg i that is operated by airline k. In this case, airline k accepts a request for itinerary j if there is enough capacity to serve a request for this itinerary and

$$p_j^k \ge \sum_{i \in \mathcal{L}^k} a_{ij} \, \hat{\pi}_i^k. \tag{21}$$

Comparing the decision rules in (5) and (21), fare allocation p_j^k essentially serves as the fare associated with itinerary j when airline k makes its booking control decisions and airline k is concerned only about the total value of the capacity consumed on the flight legs that it operates.

The appealing aspect of problem (17)-(20) is that once the fare allocations are computed, airline k can solve problem (17)-(20) by using only the capacity availability on the flight legs it operates and the expected demands for the itineraries it markets. In practice, as the sales take place and the capacities on the flight legs are depleted, it is common to resolve problem (17)-(20) periodically by replacing the right side of constraints (18) with the current remaining capacities and the right side of constraints (19) with the current remaining expected demands. Resolving problem (17)-(20) in this fashion significantly improves the performance of the policy derived from this problem. The fact that problem (17)-(20) uses only the information available to airline k becomes quite valuable when this problem is resolved periodically. In particular, airline k can resolve problem (17)-(20) as frequently as it desires without having access to the remaining capacity or remaining expected demand information from the other airlines. Once the fare allocations are computed, the airlines in the alliance are decoupled.

Two practical issues about the decision rule in (21) are worth emphasizing. As briefly mentioned at the beginning of this section, if we use the fare allocations obtained from a dual solution to problem (6)-(10), then the implicit assumption is that the airlines are willing to share their true expected demand information at least once at the beginning of the booking horizon when the fare allocations are computed. The first issue is that it may be advantageous for the airlines to falsify their expected demand information with the goal of manipulating the fare allocations to their advantage. The second issue is that our computational experiments demonstrate that the decision rule in (21) is quite effective in coordinating the decisions of the alliance and allows the alliance to obtain total expected revenues that are close to those that can be obtained by a central planner. Although the total expected revenue of the alliance is close to the total expected revenue that can be obtained by a central planner, each airline may still need to be incentivized in the correct fashion to remain true to the terms of the alliance. This is particularly the case since we work with fare allocations that are computed at the beginning of the booking horizon, whereas the capacities on the flight legs change over time as the sales take place, in which case, the fare allocations may not be aligned with the incentives of the individual airlines throughout the whole booking horizon. Anupindi et al. (2001) discuss similar coordination issues within the context of an inventory sharing model with lateral shipments between retailers. Their approach indeed provides the necessary incentives to make sure that the retailers are coordinated, but setting up their coordination mechanism also requires truthful information sharing.

In the next section, we establish a number of intuitive properties of the fare allocations that we obtain by using an optimal dual solution to problem (6)-(10). The last one of these properties provides additional support for the decision rule in (21) by showing that this decision rule is strongly related to the decision rule in (5).

7 Properties of Fare Allocations

Our goal in this section is to demonstrate that the fare allocations we obtain by using an optimal dual solution to problem (6)-(10) satisfy certain intuitive properties. Throughout this section, we use $\{\hat{\mu}_i^k : k \in \mathcal{K},\ i \in \mathcal{L}^k\}$ and $\{\hat{\alpha}_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$ to denote the optimal values of the dual variables respectively associated with constraints (7) and (9) in problem (6)-(10). In this case, we can interpret μ_i^k as the value of a unit of capacity on flight leg i operated by airline k and $\hat{\alpha}_j^k$ as the fare allocation of itinerary j to airline k. In the next proposition, we show that if an itinerary is not marketed by a particular airline, then there is a simple relationship between the fare allocation of the itinerary to this airline and the value of the capacities consumed by the itinerary. This result follows by using the dual of an equivalent reformulation of problem (6)-(10). We defer the proof to the appendix.

Proposition 1 *There exists an optimal dual solution to problem* (6)-(10) *such that if itinerary j is not marketed by airline k, then we have $\hat{\alpha}_j^k = \sum_{i \in \mathcal{L}^k} a_{ij}\, \hat{\mu}_i^k$.*

Therefore, if itinerary j is not marketed by airline k, then the fare allocation of itinerary j to airline k is exactly equal to the total value of the capacity consumed by itinerary j on the flight legs operated by airline k. In the next proposition, we give a natural implication of Proposition 1.

Proposition 2 *If itinerary j is not marketed by airline k and does not use any of the flight legs operated by airline k, then we have $\hat{\alpha}_j^k = 0$ in an optimal dual solution to problem* (6)-(10).

We observe that if itinerary j does not use any of the flight legs operated by airline k, then we have $a_{ij} = 0$ for all $i \in \mathcal{L}^k$, in which case, the proposition above follows from Proposition 1. Proposition 2 indicates that if itinerary j is not marketed by airline k and does not use any of the flight legs operated by airline k, then the fare allocation of itinerary j to airline k should be zero. It is important to use fare allocations that satisfy this property for the booking control policy derived from problem (17)-(20) to be sensible. In particular, if itinerary j is not marketed by airline k and does not use any of the flight legs operated by airline k, then the decision variable x_j^k does not appear in any of the constraints in problem (17)-(20). Therefore, the fare allocation of itinerary j to airline k has to be zero for problem (17)-(20) to have a finite optimal objective value.

In the next proposition, we give a comparison between the fare allocations of two itineraries to an airline that does not market them.

Proposition 3 *If itineraries j and j' are not marketed by airline k and they use the same flight legs operated by airline k, then we have $\hat{\alpha}_j^k = \hat{\alpha}_{j'}^k$ in an optimal dual solution to problem* (6)-(10).

The proposition above follows from Proposition 1 by noting that if itineraries j and j' use the same set of flight legs operated by airline k, then we have $a_{ij} = a_{ij'}$ for all $i \in \mathcal{L}^k$, in which case, we obtain $\hat{\alpha}_j^k = \sum_{i \in \mathcal{L}^k} a_{ij}\, \hat{\mu}_i^k = \sum_{i \in \mathcal{L}^k} a_{ij'}\, \hat{\mu}_i^k = \hat{\alpha}_{j'}^k$ by Proposition 1. Proposition 3 has an interesting implication. The fare associated with itinerary j can be substantially higher than the fare associated with itinerary j'. However, if airline k does not market any of these itineraries and these itineraries use the same set of flight legs operated by airline k, then the fare allocations of itineraries j and j' to airline k should be identical. In other words, the substantial difference in the fares associated with itineraries j and j' becomes irrelevant to airline k as long as airline k does not market these itineraries and these itineraries use the same set of flight legs operated by airline k. This simply leaves airline k indifferent between the sale of itineraries j and j'.

In the next proposition, we show that the decision rule in (21) reduces to the one in (5) under the fare allocations that we obtain by using an optimal dual solution to problem (6)-(10).

Proposition 4 *If itinerary j is marketed by airline k, then there exists an optimal dual solution to problem* (6)-(10) *such that we have $\hat{\alpha}_j^k \geq \sum_{i \in \mathcal{L}^k} a_{ij}\, \hat{\mu}_i^k$ if and only if $f_j \geq \sum_{\ell \in \mathcal{K}} \sum_{i \in \mathcal{L}^\ell} a_{ij}\, \hat{\mu}_i^\ell$.*

The proposition above follows from Proposition 1 as well. To see this, if itinerary j is marketed by airline k, then it follows from Proposition 1 that there exists an optimal dual solution to problem (6)-(10) that satisfies $\hat{\alpha}_j^\ell = \sum_{i \in \mathcal{L}^\ell} a_{ij}\, \hat{\mu}_i^\ell$ for all $\ell \in \mathcal{K} \setminus \{k\}$. In this case, if itinerary j satisfies the inequality $\hat{\alpha}_j^k \geq \sum_{i \in \mathcal{L}^k} a_{ij}\, \hat{\mu}_i^k$, then we can add the last equality for all $\ell \in \mathcal{K} \setminus \{k\}$ to this inequality to obtain $f_j = \sum_{\ell \in \mathcal{K}} \hat{\alpha}_j^\ell \geq \sum_{\ell \in \mathcal{K}} \sum_{i \in \mathcal{L}^\ell} a_{ij}\, \hat{\mu}_i^\ell$ and the desired result follows.

A useful implication of Proposition 4 is that the decision rules in (5) and (21) are equivalent to each other when we use the fare allocations $\{\hat{\alpha}_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$. In particular, under these fare allocations, problem (17)-(20) is identical to problem (13)-(16) and the latter problem is the decomposed version of problem (6)-(10). This observation implies that $\{\hat{\mu}_i^k : i \in \mathcal{L}^k\}$ are the optimal values of the dual variables associated with constraints (18) in problem (17)-(20), in which case, we can write the inequality in (21) as $\hat{\alpha}_j^k \geq \sum_{i \in \mathcal{L}^k} a_{ij}\, \hat{\mu}_i^k$. On the other hand, since problem (6)-(10) is equivalent to problem (1)-(4), if we let $\hat{\mu}_i = \hat{\mu}_i^k$ for all $k \in \mathcal{K}$ and $i \in \mathcal{L}^k$, then $\{\hat{\mu}_i : i \in \mathcal{L}\}$ are the optimal values of the dual variables associated with constraints (2) in problem (1)-(4). Therefore, we can write the inequality in (5) as $f_j \geq \sum_{\ell \in \mathcal{K}} \sum_{i \in \mathcal{L}^\ell} a_{ij}\, \hat{\mu}_i = \sum_{\ell \in \mathcal{K}} \sum_{i \in \mathcal{L}^\ell} a_{ij}\, \hat{\mu}_i^\ell$. Proposition 4 shows that the last two inequalities are equivalent, which implies that the decision rules in (5) and (21) are also equivalent.

8 Computational Experiments

In this section, we present computational experiments testing the performance of the decision rule in (21). In the process, we test the quality of the fare allocations obtained by using an optimal dual solution to problem (6)-(10). We begin by describing our benchmark methods.

8.1 Benchmark Methods

We compare the performance of three benchmark methods, corresponding to three different levels of coordination in the alliance.

Centralized Planner (CP) CP corresponds to a centralized planner that makes the decisions for the whole alliance by using problem (1)-(4). In particular, CP solves problem (1)-(4) to obtain the optimal values of the dual variables associated with constraints (2) and uses the decision rule in (5) to decide which itinerary requests to serve. In our practical implementation of CP, we divide the booking horizon into S segments and resolve problem (1)-(4) periodically at times $(s-1)\,\tau/S$ for $s = 1, \ldots, S$. To do this, noting $\int_{(s-1)\,\tau/S}^{\tau} \Lambda_j(t)\, dt$ gives the expected demand for itinerary j between the beginning of segment s and the end of the booking horizon, if the remaining capacities on the flight legs at the beginning of segment s are given by $\{c_i^s : i \in \mathcal{L}\}$, then we replace the right side of constraints (2) with $\{c_i^s : i \in \mathcal{L}\}$ and the right side of constraints (3) with $\{\int_{(s-1)\,\tau/S}^{\tau} \Lambda_j(t)\, dt : j \in \mathcal{J}\}$ and solve problem (1)-(4). We use the optimal values of the dual variables associated with constraints (2) in the decision rule in (5) until we reach the beginning of the next segment and solve problem (1)-(4) again. We use $S = 20$.

Coordinated Alliance (CA) CA corresponds to the case where we compute the fare allocations by using an optimal dual solution to problem (6)-(10), but once the fare allocations are computed, each airline makes its own booking control decisions for the itineraries that it markets. In particular, we solve problem (6)-(10) at the beginning of the booking horizon to obtain the optimal values of the dual variables associated with constraints (9). Using $\{\hat{\alpha}_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$ to denote the optimal values of the dual variables associated with these constraints, airline k uses $\{\hat{\alpha}_j^k : j \in \mathcal{J}\}$ as its fare allocations. To make its booking control decisions, airline k replaces the fare allocations $\{p_j^k : j \in \mathcal{J}\}$ in problem (17)-(20) with $\{\hat{\alpha}_j^k : j \in \mathcal{J}\}$ and solves this problem to obtain the optimal values of the dual variables associated with constraints (18). In this case, airline k uses the decision rule in (21) to decide which of the requests for the itineraries that it markets to accept. Similar to our practical implementation of CP, CA divides the booking horizon into S segments and resolves problem (17)-(20) at the beginning of each segment. Using $\{c_i^s : i \in \mathcal{L}^k\}$ to denote the remaining capacities on the flight legs operated by airline k at the beginning of segment s, airline k replaces the right side of constraints (18) with $\{c_i^s : i \in \mathcal{L}^k\}$ and the right side of constraints (19) with $\{\int_{(s-1)\,\tau/S}^{\tau} \Lambda_j(t)\, dt : j \in \mathcal{J}^k\}$ and solves problem (17)-(20). In this case, airline k uses the optimal values of the dual variables associated with constraints (18) until it reaches the beginning of the next segment and solves problem (17)-(20) again. Similar to CP, we use $S = 20$ for CA in all of our test problems.

We emphasize that when an airline resolves problem (17)-(20) periodically, it only uses the remaining capacities on the flight legs it operates and the expected demands for the itineraries it markets. In other words, airlines need to share the capacity and expected demand information only when computing the fare allocations at the beginning of the booking horizon, but not during the course of the booking horizon. If an airline markets an itinerary that uses flight legs operated by multiple airlines and the airline needs to decide whether to accept a request for this itinerary, then the marketing airline only needs to know whether there is enough capacity availability on the flight legs used by this itinerary, but not the exact amount of remaining capacity on the flight legs operated by other airlines.

Fixed Percent Fare Allocations (FP) FP corresponds to the case where the airlines use ad hoc fare allocations that are based on a fixed percentage. In particular, letting $\rho \in [0, 1]$ be the fixed percent fare allocation parameter, consider an itinerary j that uses flight legs operated by multiple airlines. In this case, if airline k_j is the marketing airline for this itinerary, then the fare allocation of itinerary j to the marketing airline k_j is simply given by $p_j^{k_j} = \rho\, f_j$, which is a fixed percentage of the total fare of the itinerary. The other airlines share the remaining portion $(1-\rho)\, f_j$ of the total fare, prorated according to the number of flight legs that they operate in this itinerary. In particular, if itinerary j is not marketed by airline k, then the fare allocation of itinerary j to airline k is given by $p_j^k = (1-\rho)\, f_j \sum_{i \in \mathcal{L}^k} a_{ij} / \sum_{k' \in \mathcal{K} \setminus \{k_j\}} \sum_{i \in \mathcal{L}^{k'}} a_{ij}$. Note that we have $\sum_{k \in \mathcal{K} \setminus \{k_j\}} p_j^k = (1-\rho)\, f_j$ so that the total fare allocation of itinerary j to all of the airlines is f_j. On the other hand, considering an itinerary j that uses flight legs operated only by the marketing airline k_j, the fare allocation of this itinerary to the marketing airline is f_j and the fare allocations to the other airlines are zero. In this way, we have a set of fare allocations $\{p_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$ for any given value of ρ. To make its booking control decisions, airline k uses the fare allocations $\{p_j^k : j \in \mathcal{J}\}$ in problem (17)-(20) and solves this problem to obtain the optimal values of the dual variables associated with constraints (18). In this case, airline

k uses the decision rule in (21) to decide which of the requests for the itineraries that it markets to accept. The performance of FP naturally depends on the value of ρ. In our computational experiments, we test the performance of FP for each value of ρ in $\{0, 0.1, 0.2, \ldots, 0.9, 1\}$ and report the performance corresponding to the best value of ρ. The best value of ρ may be different for different test problems. Similar to CP and CA, FP divides the booking horizon into S segments and resolves problem (17)-(20) at the beginning of each segment. We use $S = 20$ for FP as well.

8.2 Experimental Setup

In our computational experiments, we consider an airline network with one hub serving N spokes. There is one flight leg that connects each spoke to the hub and another one that connects the hub to each spoke. Therefore, the number of flight legs is $2N$. There is a high fare and a low fare itinerary that connect each origin destination pair, which implies that the number of itineraries is $2N\,(N+1)$. The fare associated with a high fare itinerary is κ times the fare associated with the corresponding low fare itinerary. There are K airlines in the alliance. We divide the spokes into K equal subsets and each subset of spokes is controlled by a different airline. An airline operates all of the flight legs originating or terminating at the spokes that it controls. On the other hand, an airline markets the itineraries originating at the spokes that it controls, but if an itinerary originates at the hub, then this itinerary is marketed by the airline that controls the destination spoke of the itinerary. Figure 1 shows the structure of the airline network and the alliance for the case with $N = 8$ and $K = 4$.

In all of our test problems, we normalize the total rate with which the itinerary requests arrive into the system to one. In other words, we have $\sum_{j \in \mathcal{J}} \Lambda_j(t) = 1$ for all $t \in [0, \tau]$. The length of the booking horizon is $\tau = 1{,}200$ so that the total expected number of itinerary requests over the whole booking horizon is 1,200. We calibrate the arrival rate function $\{\Lambda_j(t) : t \in [0, \tau]\}$ for the different itineraries so that low fare itineraries arrive more frequently at the beginning of the booking horizon, whereas the arrival rates of the high fare itineraries get larger later on. To do this, we generate a sample from the uniform distribution between zero and one for each origin destination pair in the airline network and normalize these samples by their sum so that they add up to one. Using P_{od} to denote this normalized sample for origin destination pair (o, d) and letting $\zeta(t)$ be an increasing function of t, we use $P_{od}\,\zeta(t)$ as the arrival rate function for the requests for the high fare itinerary that connects origin destination pair (o, d). Similarly, we use $P_{od}\,(1 - \zeta(t))$ as the arrival rate function for the requests for the low fare itinerary that connects origin destination pair (o, d). In our computational experiments, we have $\zeta(t) = 0$ until one third of the booking horizon. After this point, $\zeta(t)$ increases linearly to one.

Noting that the expected demand for the capacity on flight leg i is given by $\sum_{j \in \mathcal{J}} a_{ij}\,\lambda_j$, we measure the tightness of the leg capacities by $\gamma = \frac{\sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} a_{ij}\,\lambda_j}{\sum_{i \in \mathcal{L}} c_i}$. We label our test problems by using the tuple (N, K, κ, γ) and vary (N, K, κ, γ) over $\{8, 16\} \times \{2, 4, 8\} \times \{4, 6\} \times \{1.0, 1.3, 1.6\}$. In this way, we obtain 36 test problems for our computational experiments. With this experimental setup, the capacities on the flight legs range between 30 and 180 depending on the number of flight legs in the airline network and the tightness of the leg capacities.

8.3 Computational Results

Table 1 summarizes our main computational results. The left portion of Table 1 focuses on the test problems with eight spokes in the airline network, whereas the right portion focuses on the test problems with 16 spokes. The first column in the table shows the characteristics of the test problems by using the tuple (N, K, κ, γ). The second, third and fourth columns respectively show the total expected revenues obtained by CP, CA and FP for the whole airline alliance. These total expected revenues are estimated by simulating the performance of the benchmark methods under common random numbers. The fifth column shows the percent gaps between the total expected revenues obtained by CP and CA. We also include a checkmark (✓) in this column if the percent gap between the total expected revenues obtained by CP and CA is statistically significant at 95% level and a dash ($-$) if there is no statistically significant difference. The sixth column does the same thing as the fifth column, but it focuses on the percent gaps between the total expected revenues obtained by CP and FP. CP corresponds to a centralized planner making the booking control decisions for the whole alliance, whereas CA and FP allow the airlines to make their booking control decisions autonomously. Our goal in comparing the total expected revenues obtained by CA and FP with the ones obtained by CP is that we would like to see how much revenue loss we incur by allowing the airlines to make their booking control decisions autonomously.

The results in Table 1 indicate that CA does a good job of coordinating the decisions of the airlines, whereas FP can perform poorly. Over the test problems with eight spokes, the average performance gap between CP and CA is 2.81%. The same average performance gap comes out to be 2.71% over the test problems with 16 spokes. In other words, if the airlines choose the fare allocations by solving problem (6)-(10) once at the beginning of the booking horizon, but otherwise make their booking control decisions autonomously, then the total expected revenues obtained by the alliance deviate from the ones obtained by a central planner by 2% to 3% on average. There are five test problems where the total expected revenues obtained by CA and CP deviate from each other by more than 5%. Overall, the revenue loss of CA incurred by making autonomous booking control decisions is quite reasonable, especially considering the fact that CA does not require the airlines to share remaining capacity or remaining expected demand information with each other over the booking horizon. For many test problems the revenue loss of CA is on the order of 1% to 4%. There are three test problems, where the revenue loss is on the order of 6% to 7% and we revisit these test problems later in this section.

The ad hoc fare allocations used by FP may not perform well. Over the test problems with eight spokes, the average performance gap between CP and FP is 12.90%, whereas the same average performance gap over the test problems with 16 spokes is 13.52%. The smallest performance gap between CP and FP is on the order of 6% and the performance gap between CP and FP exceeds 10% in 27 out of 36 test problems. The results indicate that choosing the fare allocations carefully is crucial to coordinate the decisions of the airlines and to come up with an efficient booking control mechanism that obtains total expected revenues close to those obtained by a central planner.

CA computes the fare allocations once at the beginning and uses these fare allocations throughout the booking horizon. Although the airlines in the alliance resolve problem (17)-(20) periodically to adjust their booking control policies, the fare allocations used in this problem do not change. An interesting question is how much the performance of CA improves when the fare allocations are periodically readjusted based on the remaining capacities on the flight legs and remaining expected demands to come. To answer this question, we divide the booking horizon into L segments and recompute the fare allocations at times $(l-1)\tau/L$ for $l = 1, \ldots, L$. In particular, if the remaining capacities on the flight legs at the beginning of segment l is given by $\{c_i^l : i \in \mathcal{L}\}$, then we replace the right side of constraints (7) with $\{c_i^l : k \in \mathcal{K},\ i \in \mathcal{L}^k\}$ and the right side of constraints (8) with $\{\int_{(l-1)\tau/L}^{\tau} \Lambda_j(t)\, dt : k \in \mathcal{K},\ j \in \mathcal{J}^k\}$ and solve problem (6)-(10). We use the optimal values of the dual variables associated with constraints (9) as the fare allocations until we reach the beginning of the next segment. Solving this problem requires the airlines to get together and share remaining capacity and remaining expected demand information with each other. Therefore, it is sensible to choose a small value for L. We choose $L = 2$ in our computational experiments, which corresponds to recomputing the fare allocations only at the beginning and at the middle of the booking horizon.

Table 2 compares the total expected revenues obtained by CP and CA when we recompute the fare allocations for CA twice over the booking horizon. The second column in this table shows the total expected revenues obtained by CP and it is identical to the second column in Table 1. The third column in Table 2 shows the total expected revenues obtained by CA when we recompute the fare allocations twice over the booking horizon. The fourth column shows the percent gaps between the total expected revenues in the second and third columns. Comparing the performance gaps between CP and CA in Table 2 with those in Table 1, we observe that recomputing the fare allocations just twice makes a noticeable impact on the performance of CA. In Table 2, over the test problems with eight and 16 spokes, the average performance gaps between CP and CA are respectively 1.82% and 1.49%, whereas the same average performance gaps are respectively 2.81% and 2.71% in Table 1. We recall that there is a test problem in Table 1, where the total expected revenue obtained by CP exceeds the one obtained by CA by 7.19%. The performance gap between CP and CA for this test problem goes down to 2.76% when we recompute the fare allocations twice over the booking horizon. Similar observations can be made for the other two test problems in Table 1 with more than 6% performance gap between CP and CA. Therefore, while CA generally does a good job of coordinating the decisions of the airlines even when the fare allocations are computed once at the beginning, there can be noticeable value in recomputing the fare allocations as little as twice over the booking horizon. Finally, it is also worthwhile to point out that there is one test problem in Table 2, where CA performs better than CP with a statistically significant performance gap. If the centralized planner made the optimal booking control decisions, then the total expected revenues obtained by CP would always exceed those obtained by CA. However, given that the booking control policies used by CP and CA are both approximate, it is indeed possible that CA performs better than CP.

9 Conclusions

In this paper, we developed a revenue management model for airline alliances. The fundamental idea behind our model is to start with a monolithic deterministic linear programming approximation that models the decisions made by a central planner for the whole alliance. By associating dual multipliers with the constraints that link the decisions for different airlines, we relax these constraints and decompose the linear program by the airlines. In this case, we can extract a booking control policy from the relaxed linear program that allows each airline to make its own capacity allocation decisions autonomously. Furthermore, the dual multipliers indicate how the revenues from different itineraries should be shared among the airlines. A natural direction for further work is that there are numerous refinements to the deterministic linear program aimed at capturing the stochastic and dynamic nature of the itinerary requests more accurately. One can follow the same approach in this paper to use these refined linear programs in the alliance setting. It would be interesting and practically useful to see what benefits can be obtained from the refined linear programs. In addition, as we mention in the paper, computing the fare allocations by using our approach requires the airlines to truthfully share expected demand and capacity information at least once at the beginning of the booking horizon when the fare allocations are computed, but this information may be falsified to change the fare allocations to the advantage of a particular alliance member. Thus, it would be useful to study airline alliance models that provide incentives for truthful information sharing.

References

Adelman, D. (2007), 'Dynamic bid-prices in revenue management', *Operations Research* **55**(4), 647–661.

Anupindi, R., Bassok, Y. and Zemel, E. (2001), 'A general framework for the study of decentralized distribution systems', *Manufacturing & Service Operations Management* **3**(4), 349–368.

Boyd, E. A. (1998), Airline alliance revenue management, Technical report, Pros Strategic Solutions, Houston, TX.

Dantzig, G. and Wolfe, P. (1960), 'Decomposition principle for linear programs', *Operations Research* **8**, 101–111.

Erdelyi, A. and Topaloglu, H. (2011), 'Using decomposition methods to solve pricing problems in network revenue management', *Journal of Revenue and Pricing Management* **10**(4), 325–343.

Houghtalen, L., Ergun, O. and Sokol, J. (2007), Designing mechanisms for the management of carrier alliances, Technical report, Georgia Institute of Technology, Atlanta, GA.

Hu, X., Caldentey, R. and Vulcano, G. (2010), Revenue sharing in airline alliances, Technical report, New York University, New York, NY.

Liu, Q. and van Ryzin, G. (2008), 'On the choice-based linear programming model for network revenue management', *Manufacturing & Service Operations Management* **10**(2), 288–310.

Simpson, R. W. (1989), Using network flow techniques to find shadow prices for market and seat inventory control, Technical report, Massachusetts Institute of Technology Flight Transportation Laboratory Memorandum M89-1, Cambridge, MA.

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. (1999), 'A randomized linear programming method for computing network bid prices', *Transportation Science* **33**(2), 207–216.

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* **4**(1), 66–82.

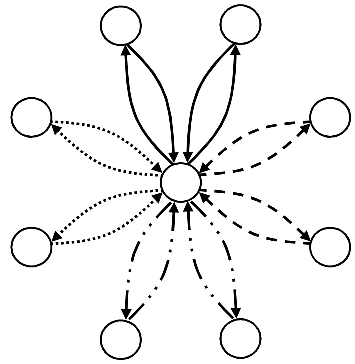

Figure 1: Structure of the airline network and the flight legs operated by the airlines with $N = 8$ and $K = 4$. The arcs in different styles represent the flight legs operated by different airlines.

Williamson, E. L. (1992), Airline Network Seat Control, PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, Cambridge, MA.

Wright, C. P. (2011), Decoupling airline alliances: A bid-price approach to the game of incomplete information, Technical report, Niagara University, Lewiston, NY.

Wright, C. P., Groenevelt, H. and Shumsky, R. A. (2010), 'Dynamic revenue management in airline alliances', *Transportation Science* **44**(1), 15–37.

Zhang, D. (2011), 'An improved dynamic programming decomposition approach for network revenue management', *Manufacturing & Service Operations Management* **13**(1), 35–52.

Problem	Tot. Exp. Rev.			Comp. with CP	
(N, K, κ, γ)	CP	CA	FP	CA	FP
$(8, 2, 4.0, 1.0)$	191,624	185,861	176,766	3.01 (✓)	7.75 (✓)
$(8, 2, 4.0, 1.3)$	178,871	177,529	166,370	0.75 (✓)	6.99 (✓)
$(8, 2, 4.0, 1.6)$	168,862	166,976	157,090	1.12 (✓)	6.97 (✓)
$(8, 2, 6.0, 1.0)$	255,923	251,286	234,305	1.81 (✓)	8.45 (✓)
$(8, 2, 6.0, 1.3)$	242,819	241,755	223,872	0.44 (✓)	7.80 (✓)
$(8, 2, 6.0, 1.6)$	232,426	228,133	214,976	1.85 (✓)	7.51 (✓)
$(8, 4, 4.0, 1.0)$	191,624	180,593	164,394	5.76 (✓)	14.21 (✓)
$(8, 4, 4.0, 1.3)$	178,871	172,508	152,064	3.56 (✓)	14.99 (✓)
$(8, 4, 4.0, 1.6)$	168,862	164,474	144,166	2.60 (✓)	14.63 (✓)
$(8, 4, 6.0, 1.0)$	255,923	247,138	217,891	3.43 (✓)	14.86 (✓)
$(8, 4, 6.0, 1.3)$	242,819	239,046	201,659	1.55 (✓)	16.95 (✓)
$(8, 4, 6.0, 1.6)$	232,426	228,741	199,172	1.59 (✓)	14.31 (✓)
$(8, 8, 4.0, 1.0)$	191,624	177,841	162,388	7.19 (✓)	15.26 (✓)
$(8, 8, 4.0, 1.3)$	178,871	169,970	148,479	4.98 (✓)	16.99 (✓)
$(8, 8, 4.0, 1.6)$	168,862	163,413	139,788	3.23 (✓)	17.22 (✓)
$(8, 8, 6.0, 1.0)$	255,923	244,278	215,234	4.55 (✓)	15.90 (✓)
$(8, 8, 6.0, 1.3)$	242,819	237,329	201,636	2.26 (✓)	16.96 (✓)
$(8, 8, 6.0, 1.6)$	232,426	230,373	199,018	0.88 (✓)	14.37 (✓)
Average				2.81	12.90

Problem	Tot. Exp. Rev.			Comp. with CP	
(N, K, κ, γ)	CP	CA	FP	CA	FP
$(16, 2, 4.0, 1.0)$	182,241	172,017	161,137	5.61 (✓)	11.58 (✓)
$(16, 2, 4.0, 1.3)$	168,033	164,313	148,120	2.21 (✓)	11.85 (✓)
$(16, 2, 4.0, 1.6)$	157,633	156,408	138,640	0.78 (✓)	12.05 (✓)
$(16, 2, 6.0, 1.0)$	245,166	237,440	211,190	3.15 (✓)	13.86 (✓)
$(16, 2, 6.0, 1.3)$	230,285	229,369	204,747	0.40 (✓)	11.09 (✓)
$(16, 2, 6.0, 1.6)$	219,363	219,615	198,633	-0.11 (–)	9.45 (✓)
$(16, 4, 4.0, 1.0)$	182,241	169,925	150,908	6.76 (✓)	17.19 (✓)
$(16, 4, 4.0, 1.3)$	168,033	162,175	139,088	3.49 (✓)	17.23 (✓)
$(16, 4, 4.0, 1.6)$	157,633	154,947	134,392	1.70 (✓)	14.74 (✓)
$(16, 4, 6.0, 1.0)$	245,166	235,546	205,596	3.92 (✓)	16.14 (✓)
$(16, 4, 6.0, 1.3)$	230,285	227,491	204,614	1.21 (✓)	11.15 (✓)
$(16, 4, 6.0, 1.6)$	219,363	218,791	198,679	0.26 (✓)	9.43 (✓)
$(16, 8, 4.0, 1.0)$	182,241	169,154	147,947	7.18 (✓)	18.82 (✓)
$(16, 8, 4.0, 1.3)$	168,033	161,439	137,886	3.92 (✓)	17.94 (✓)
$(16, 8, 4.0, 1.6)$	157,633	154,325	133,580	2.10 (✓)	15.26 (✓)
$(16, 8, 6.0, 1.0)$	245,166	234,998	208,835	4.15 (✓)	14.82 (✓)
$(16, 8, 6.0, 1.3)$	230,285	226,801	204,520	1.51 (✓)	11.19 (✓)
$(16, 8, 6.0, 1.6)$	219,363	218,283	198,327	0.49 (✓)	9.59 (✓)
Average				2.71	13.52

Table 1: Total expected revenues obtained by CP, CA and FP.

Problem (N, K, κ, γ)	Tot. Exp. Rev. CP	CA	Comp. betw. CP-CA
$(8, 2, 4.0, 1.0)$	191,624	187,030	2.40 (✓)
$(8, 2, 4.0, 1.3)$	178,871	177,551	0.74 (✓)
$(8, 2, 4.0, 1.6)$	168,862	167,065	1.06 (✓)
$(8, 2, 6.0, 1.0)$	255,923	251,395	1.77 (✓)
$(8, 2, 6.0, 1.3)$	242,819	241,591	0.51 (✓)
$(8, 2, 6.0, 1.6)$	232,426	227,859	1.96 (✓)
$(8, 4, 4.0, 1.0)$	191,624	187,400	2.20 (✓)
$(8, 4, 4.0, 1.3)$	178,871	173,964	2.74 (✓)
$(8, 4, 4.0, 1.6)$	168,862	164,639	2.50 (✓)
$(8, 4, 6.0, 1.0)$	255,923	251,976	1.54 (✓)
$(8, 4, 6.0, 1.3)$	242,819	240,774	0.84 (✓)
$(8, 4, 6.0, 1.6)$	232,426	228,301	1.77 (✓)
$(8, 8, 4.0, 1.0)$	191,624	186,334	2.76 (✓)
$(8, 8, 4.0, 1.3)$	178,871	173,857	2.80 (✓)
$(8, 8, 4.0, 1.6)$	168,862	164,766	2.43 (✓)
$(8, 8, 6.0, 1.0)$	255,923	251,074	1.89 (✓)
$(8, 8, 6.0, 1.3)$	242,819	237,370	2.24 (✓)
$(8, 8, 6.0, 1.6)$	232,426	230,955	0.63 (✓)
Average			1.82

Problem (N, K, κ, γ)	Tot. Exp. Rev. CP	CA	Comp. betw. CP-CA
$(16, 2, 4.0, 1.0)$	182,241	177,058	2.84 (✓)
$(16, 2, 4.0, 1.3)$	168,033	165,415	1.56 (✓)
$(16, 2, 4.0, 1.6)$	157,633	156,610	0.65 (✓)
$(16, 2, 6.0, 1.0)$	245,166	240,013	2.10 (✓)
$(16, 2, 6.0, 1.3)$	230,285	229,426	0.37 (✓)
$(16, 2, 6.0, 1.6)$	219,363	219,901	-0.25 (✓)
$(16, 4, 4.0, 1.0)$	182,241	177,366	2.67 (✓)
$(16, 4, 4.0, 1.3)$	168,033	163,916	2.45 (✓)
$(16, 4, 4.0, 1.6)$	157,633	155,455	1.38 (✓)
$(16, 4, 6.0, 1.0)$	245,166	240,991	1.70 (✓)
$(16, 4, 6.0, 1.3)$	230,285	227,882	1.04 (✓)
$(16, 4, 6.0, 1.6)$	219,363	219,395	-0.01 (−)
$(16, 8, 4.0, 1.0)$	182,241	177,274	2.73 (✓)
$(16, 8, 4.0, 1.3)$	168,033	163,410	2.75 (✓)
$(16, 8, 4.0, 1.6)$	157,633	154,871	1.75 (✓)
$(16, 8, 6.0, 1.0)$	245,166	241,002	1.70 (✓)
$(16, 8, 6.0, 1.3)$	230,285	227,558	1.18 (✓)
$(16, 8, 6.0, 1.6)$	219,363	218,955	0.19 (−)
Average			1.49

Table 2: Total expected revenues obtained by CP and CA when we recompute the fare allocations for CA twice over the booking horizon.

Appendix: Proof of Proposition 1

The key observation is that we only need to impose nonnegativity constraints on the decision variables $\{x_j^k : k \in \mathcal{K},\ j \in \mathcal{J}^k\}$ in problem (6)-(10), in which case, problem (6)-(10) equivalently becomes

$$
\begin{aligned}
\max \quad & \sum_{j \in \mathcal{J}} f_j\, x_j^{\psi} && && (22)\\
\text{subject to} \quad & \sum_{j \in \mathcal{J}} a_{ij}\, x_j^k \le c_i && \forall\, k \in \mathcal{K},\ i \in \mathcal{L}^k && (23)\\
& x_j^k \le \lambda_j && \forall\, k \in \mathcal{K},\ j \in \mathcal{J}^k && (24)\\
& x_j^{\psi} - x_j^k = 0 && \forall\, k \in \mathcal{K},\ j \in \mathcal{J} && (25)\\
& x_j^k \ge 0 && \forall\, k \in \mathcal{K},\ j \in \mathcal{J}^k. && (26)
\end{aligned}
$$

To see the equivalence between problems (6)-(10) and (22)-(26), we choose a decision variable x_j^k such that $j \notin \mathcal{J}^k$. This decision variable has an explicit nonnegativity constraint in problem (6)-(10), but not in problem (22)-(26). We proceed to show that this decision variable x_j^k always takes a nonnegative value in a feasible solution to problem (22)-(26). We use k' to denote the airline that markets itinerary j so that $j \in \mathcal{J}^{k'}$. By constraints (25), we have $x_j^k = x_j^{\psi} = x_j^{k'} \ge 0$, where the last inequality follows from constraints (26) and the fact that $j \in \mathcal{J}^{k'}$. This shows that the decision variable x_j^k always takes a nonnegative value in a feasible solution to problem (22)-(26). Therefore, all of the decision variables $\{x_j^k : k \in \mathcal{K},\ j \in \mathcal{J} \setminus \mathcal{J}^k\}$ take nonnegative values in a feasible solution to problem (22)-(26), which implies that problems (6)-(10) and (22)-(26) are equivalent to each other.

To complete the proof of Proposition 1, we observe that if airline k does not market itinerary j, then we have $j \notin \mathcal{J}^k$ and the decision variable x_j^k appears only in constraints (23) and (25) in problem (22)-(26) with no explicit nonnegativity constraint. We use $\{\hat{\mu}_i^k : k \in \mathcal{K},\ i \in \mathcal{L}^k\}$ and $\{\hat{\alpha}_j^k : k \in \mathcal{K},\ j \in \mathcal{J}\}$ to respectively denote the optimal values of the dual variables associated with constraints (23) and (25). In this case, the constraint associated with the decision variable x_j^k in the dual of problem (22)-(26) implies that $\sum_{i \in \mathcal{L}^k} a_{ij}\, \hat{\mu}_i^k - \hat{\alpha}_j^k = 0$, which is the desired result.