

Dynamic Revenue Management in Airline Alliances

Christopher P. Wright

College of Business Administration, Niagara University, Niagara, New York 14109,
cwright@niagara.edu

Harry Groenevelt

Simon Graduate School of Business, University of Rochester, Rochester, New York 14627,
groenevelt@simon.rochester.edu

Robert A. Shumsky

Tuck School of Business, Dartmouth College, Hanover, New Hampshire 03755,
robert.shumsky@dartmouth.edu

Major airlines are selling increasing numbers of interline itineraries in which flights operated by two or more airlines are combined and sold together. One reason for this increase is the rapid growth of airline alliances, which promote the purchase of interline itineraries and, therefore, virtually extend the reach of each alliance member's network. This practice, however, creates a difficult coordination problem: Each member of the alliance makes revenue management decisions to maximize its own revenue and the resulting behavior may produce suboptimal revenue for the alliance as a whole. Airline industry researchers and consultants have proposed a variety of static and dynamic mechanisms to control revenue management decisions across alliances (a dynamic mechanism adjusts its parameters as the number of available seats in the network changes and time passes). In this paper, we formulate a Markov game model of a two-partner alliance that can be used to analyze the effects of these mechanisms on each partner's behavior. We begin by showing that no Markovian transfer pricing mechanism can coordinate an arbitrary alliance. Next, we examine three dynamic schemes as well as three forms of the static scheme widely used in practice. We derive the equilibrium acceptance policies under each scheme and use analytical techniques as well as numerical analyses of sample alliances to generate fundamental insights about partner behavior under each scheme. The analysis and numerical examples also illustrate how certain transfer price schemes are likely to perform in networks with particular characteristics.

Key words: revenue management; yield management; airlines; alliances; Markov game

History: Received: February 2008; revision received: June 2009; accepted: July 2009.

1. Introduction

When one of the authors recently planned a trip from Boston to Barcelona, British Airways offered a convenient itinerary for \$823. The itinerary began with a leg on British Airways from Boston to London followed by a second leg on another airline, Iberia, to Barcelona. We will call such an itinerary an *interline itinerary* because it includes service on multiple airlines. The availability of this two-leg interline itinerary for this fare is contingent on two decisions: (1) the airline that sells the ticket (the *marketing airline*, in this example, British Airways) must make a seat available on one leg, and (2) the operator of the other leg (the *operating airline*, in this example, Iberia) must agree to accept the connecting passenger from the marketing airline. Because both airlines practice revenue management, these decisions depend on the price paid by the consumer to the marketing airline for the ticket *and* the price paid by the marketing air-

line to the operating airline for the use of a seat. This article examines agreements among airlines that govern the latter price, sometimes called revenue sharing, transfer price, or proration agreements. We show how these agreements have subtle and potentially significant effects on individual airline behavior as well as on the total revenue collected by multiple airlines across their combined networks.

It is becoming increasingly important to understand the impact of proration agreements on airline revenue management because sales of interline itineraries have been growing. This increase is due in part to a type of marketing arrangement called a code share agreement. Under this arrangement, the operating airline's flight is also listed as a flight with the marketing airline's name. In the example above, the second leg had an Iberia flight number but was also labeled British Airways Flight 7073 from London to Barcelona. Analysis of data from the U.S. Department of Transportation (2006) reveals that the fraction

of interline itineraries within the United States rose from 10% in 1998 to 20% in 2004, and most of those interline itineraries were marketed under code share agreements. Overall, 46% of revenues collected from U.S. domestic flights in 2004 came from interline itineraries.

A second factor driving up interline traffic is the growth of airline alliances. These alliances usually combine code share agreements with other arrangements such as schedule coordination and the merger of frequent flier programs. In March 2006, 59% of all worldwide ASMs (available seat miles, a measure of total capacity equal to the number of seats multiplied by the number of miles flown and summed over all flights) were flown by airlines belonging to one of the three largest international alliances—Star, SkyTeam, or oneworld (Lott 2006). Both international and U.S. alliance activity is expected to continue growing (Lott 2006, Belden 2007, Shumsky 2006).

These international alliances facilitate formal marketing and operational arrangements among airlines, but this paper uses the term alliance in a much weaker sense: an alliance is formed by any two airlines that exchange interline passengers and that have a proration agreement for the revenue collected from the sale of interline itineraries.

In practice, the rules for revenue sharing are usually laid out in special prorate agreements (SPAs) that are negotiated by alliance partners. In the absence of an SPA, airlines follow the rules set out by the International Air Transport Association (IATA) (see IATA 2007 for details of the rules to be implemented over the next few years). Whether they are encoded in an SPA or by the IATA, most rules include fixed transfer prices for particular flight and fare-class combinations or other simple allocation procedures such as a split in revenue (a *proration rate*) based on relative mileage. For example, under such a mileage proration scheme, British Airways would receive the lion's share of the \$823 in the earlier example for its Boston-to-London flight. Throughout this paper, we call such rules *static schemes* because for they do not adjust proration rates as time passes, demand is realized, and seats are sold.

Although static schemes are easy to manage, they can lead to suboptimal decision making by member airlines and lost revenue for the alliance as a whole. For our example, under mileage proration, Iberia would receive a relatively small share of the ticket revenue when flying interline passengers. Therefore, Iberia may choose to focus on its own *intra-line* (those taking itineraries contained entirely on Iberia's own network) customers and not hold seats for British Airways customers who may be more lucrative for the alliance. The underlying flaw in any static proration scheme is that the revenue-sharing proportions are

not adjusted to reflect the actual value of seat inventory. Because the revenue management system of each airline in the alliance maximizes the revenue of that airline, an airline may reject an itinerary if the transfer price undervalues the real-time value of its seats, even if the total revenue from the itinerary is large.

Given the deficiencies of static schemes, major airlines are considering dynamic schemes such as the use of the real-time opportunity costs of seats (or "bid prices") as transfer prices. In this paper, we examine dynamic schemes that have been described in the published literature (see Vinod 2005) and have been suggested to the authors by industry executives and revenue managers. In the industry, there is interest in dynamic schemes but also much uncertainty. There are certainly technical and legal barriers to implementation—for example, antitrust legislation in the United States prohibits the exchange of certain types of information among airlines—but another significant barrier is uncertainty over how revenue-maximizing airlines would respond to dynamic schemes and whether such schemes would produce real benefits to the alliance. In fact, there is no published literature on this topic and, to our knowledge there has been no rigorous analysis of these effects by researchers within the industry. This paper is a first attempt to fill that gap.

Now, we summarize the organization of the paper and its results. After reviewing the relevant literature in §2, in §3 we describe our general model for a two-airline alliance. In §4, we describe a centralized network and determine the first-best policies for a centralized yield management system. In §5, we describe a two-airline alliance model and analyze static and dynamic transfer price schemes. In this section, we first use a counterexample to show that no dynamic scheme is guaranteed to maximize alliancewide revenue, unless the dynamic scheme includes revenue-sharing rules that depend on the sample path of inventory sales (note that a scheme based on sample paths would be orders of magnitude more complex than the static and dynamic schemes being considered by the airlines). We then derive equilibrium policies for the alliance partners under certain dynamic schemes and we use the analysis to highlight the strengths and weaknesses of the schemes in terms of total alliance revenue. In §6, we describe numerical experiments that support the insights from §5. The experiments also compare the performance of static and dynamic schemes given certain network parameters. We find that static schemes can perform as well as dynamic schemes for certain networks but that the performance of a static scheme that is optimal for one network can degrade quickly as the network parameters change. Dynamic schemes often perform better and are more robust. We find,

however, that the performance of dynamic schemes can be significantly reduced if each operating airline chooses a transfer price to maximize its own revenue. Such would be the case, for example, if the partners initially agree to use bid prices as transfer prices but then each partner attempts to increase its revenue by reporting incorrect bid prices. Finally, in §7, we summarize our results and describe future research.

There are a few caveats for the results in this paper. First, we make strong assumptions about the amount of information available to each alliance partner. Specifically, we use a Markov game model to describe the alliance and use the Nash equilibrium to describe the airlines' behavior under each proration scheme. For legal and technical reasons, the airlines cannot coordinate their revenue management systems so it is appropriate to use the tools of noncooperative game theory. To keep the problem tractable, however, we assume that the airlines share the same information about the state of the game and the distribution of future events, e.g., each airline has perfect information about its partner's inventory level and both have identical forecasts of future arrival probabilities and revenue distributions over the entire alliance network (in technical terms, we define a game of complete information). Although this assumption is not realistic, we believe that the amount of transparency in the industry is increasing. For example, airlines regularly use the Web to monitor the lowest available fare of their competitors and hire "market intelligence services" such as QL2 ([www.ql2.com](http://www ql2.com)) to gather information about competitor actions. Our model represents a logical extreme case and the full-information assumption allows us to generate fundamental insights on how certain proration schemes behave. We provide additional details and discussion of our information-sharing assumptions in §3.2. In general, analysis of games with incomplete information will be an interesting area for further research.

A second caveat is that the numerical experiments described in §6 were conducted using small networks in terms of the number of flights and the number of seats. Again, our purpose is to gain basic insights, i.e., to identify the fundamental advantages and disadvantages of each transfer price scheme. In addition, we demonstrate that many of these insights apply as the number of seats in the network grows. An important area for additional research, however, will be to examine alliance performance over networks of realistic size.

Finally, at a higher level than our analysis, the alliance partners are engaged in a cooperative process to determine which routes should be available for interline traffic and what proration rules to use on those routes. In general, the partners seek to increase alliancewide revenue and to allocate revenues so

that all members are willing to participate in the alliance. We do not model this higher-level process. Instead, our model provides information about how airlines behave and how total revenues are affected given sets of interline routes and particular proration rules. Models that focus on the higher-level problem have begun to appear in the literature. For example, Agarwal and Ergun (2007) examine the allocation mechanism design problem for cargo shippers. For airline revenue management, successful high-level negotiations depend on information about the effects of particular proration schemes on network revenues. To our knowledge, the model presented here is the first to provide such information.

2. Literature Review

Revenue management (also referred to as yield management) and its application to the airline industry have received a great deal of attention since the 1970s when Littlewood (1972) first described the basic problem. In that article, Littlewood introduces the result (now referred to as Littlewood's rule) that a request for a seat should be fulfilled only if its revenue exceeds the expected future value of the seat in question. This intuitive rule forms the basis of many control policies in both theory and practice.

Numerous authors have expanded on Littlewood's work. See, for example, Belobaba (1989), who examines a problem with multiple fare-restriction combinations; Glover et al. (1982), who look at the passenger mix problem in a network environment; You (1999), who examines a dynamic pricing model; and Talluri and van Ryzin (2004a), who use a discrete choice model of demand. For a more thorough description of the revenue management literature, see the survey by McGill and van Ryzin (1999) and the book by Talluri and van Ryzin (2004b).

The use of competitive game theory in revenue management has been limited. Vulcano, van Ryzin, and Maglaras (2002) examine a dynamic game in which a seller faces a sequence of customers who compete with each other in an auction for a fixed number of units. Netessine and Shumsky (2004) examine both horizontal and vertical competition between two airlines where each airline flies a single leg.

Several aspects of airline alliances have been examined in the literature. Barron (1997) discusses many of the legal implications of airline alliances, focusing on code-sharing agreements used widely in the industry. Park (1997) and Brueckner (2001) examine the economic effects of alliances on fares, traffic levels, profits, and market welfare. Brueckner and Whalen (2000) provide an empirical analysis of the effects of international alliances on fares, showing that interline fares charged by alliances are approximately 25% lower

than those charged by nonallied carriers. Ito and Lee (2007) examine the impact of domestic alliances on airfares.

Little attention, however, has been given to how revenue management should be implemented by an airline alliance. Wynne (1995) describes simple transfer price schemes based on the value of local fares. Boyd (1998a) discusses the methodological and technical challenges of the alliance revenue management problem. He also refers to a more formal analysis in an unpublished working paper (Boyd 1998b) in which he formulates a static linear program to describe the alliance revenue management problem. Boyd (1998b) then derives conditions under which the seat allocation between the two airlines maximizes alliancewide revenue under this model. Vinod (2005) describes many of the alliance coordination mechanisms now being considered by the airlines but provides no formal analysis of their advantages and disadvantages. Some of the schemes analyzed in our paper correspond to mechanisms described by Vinod (2005). Shumsky (2006) argues that low-cost competitors are driving the network airlines to rely on alliances for an increasing proportion of their traffic. Both Shumsky (2006) and Fernandez de la Torre (1999) discuss the need for more research on the effectiveness of alliance agreements, a need we attempt to fill here. In their paper on revenue management games, Netessine and Shumsky (2004) describe and analyze a static alliance revenue-sharing mechanism for a two-leg network based on the expected flow of passengers. In our paper, we analyze the performance of dynamic coordination mechanisms that are designed for arbitrary alliance networks and are similar to schemes that are proposed by, or actually used by, the airlines.

Ongoing research by Houghtalen, Ergun, and Sokol (2007) and Agarwal and Ergun (2007) looks at various aspects of alliances, focusing specifically on cargo carriers. In addition to the inherent difference between the cargo and passenger revenue management problems (see Kasilingam 1996), their analysis differs from ours in two fundamental ways: (1) they focus on the high-level alliance formation problem with cooperative game theory as the appropriate method, while we formulate a noncooperative game given an existing alliance and particular revenue-sharing rules; and (2) they focus on a deterministic optimization problem in which all demand for cargo service has been realized before routing decisions are made, while our passenger yield management problem is most appropriately described by a model in which demand is uncertain and arrives over time.

Finally, our paper is related to the extensive literature on supply chain coordination (see Nagarajan and Sošić (2008) and Cachon and Netessine (2004) for

overviews of the related literature that use, respectively, cooperative and competitive game theory). There are several attributes of our problem, however, that distinguish it from this research stream. First, the flow of products in the traditional supply chain literature moves in one direction—from raw materials to the consumer. Therefore, the unused product does not move “sideways” within a level. Second, in the supply chain literature, production of a product begins at one level with one set of firms (suppliers) and demand is fulfilled at another level by another set of firms (retailers). Neither attribute holds for our problem. For a specific contrast, consider the literature on assembly systems—we can think of a multileg itinerary as a final product assembled from multiple components. In the traditional supply chain literature, an assembler receives components from several suppliers and combines them to create a new product to sell (e.g., Nagarajan and Bassok 2008 and Granot and Yin 2008.) In our model, *either* airline may serve as the marketing airline (the de facto assembler) and either airline may serve as the operating airline (the de facto supplier).

In addition, the traditional research on supply chain coordination focuses on either single-period newsvendor problems (e.g., Lariviere and Porteus 2001) or repeated games in which inventory is replenished between each repetition of the game (e.g., Cachon and Zipkin 1999). The characteristics of such problems are quite different from ours because we study a finite, multiperiod problem with fixed capacity allocated to a stochastic arrival stream. Certain results from our paper may be similar in interpretation to results from the research on the economics of supply chains. For example, the effect of the partner price scheme in §5.4.3 can be seen as a form of double marginalization (Spengler 1950). In general, however, our problem context, model, and key results are quite different from those in the supply chain literature.

3. General Alliance Network Model

We consider a dynamic model of an alliance consisting of two partner airlines (carriers) indexed by $c \in \{1, 2\}$ (in a slight abuse of notation, we will denote the “other” airline by $-c$ instead of by $3 - c$). Our model can be seen as an extension of the network model described by Talluri and van Ryzin (1998) into a two-player game framework.

Each flight leg in the network is characterized by an origin, a destination, and a departure time (for the remainder of this paper, the terms “flight” and “flight leg” are used interchangeably). The number of flights operated by airline c is denoted m_c , and $m \equiv m_1 + m_2$ is the total number of flights offered by the alliance. The alliance offers n itineraries and each itinerary is either

a single flight or a series of connecting flights within one or both networks. The set of all itineraries is denoted N and has cardinality n . Within the alliance, these itineraries are divided into three subsets: those that involve only airline 1's flights (N_1), those that involve only airline 2's flights (N_2), and those that use flights from both airlines (N_S). Let n_1 , n_2 , and n_S be the cardinality of each subset so that $n = n_1 + n_2 + n_S$. We will refer to the sets N_1 and N_2 as *intra-line itineraries* because all legs within each set are operated by the same airline, and will refer to the set N_S as *interline itineraries* because at least one leg on any itinerary in N_S will not be operated by the airline that sold the ticket.

We use the matrix A to specify the inventory requirements of the itineraries offered by the alliance. The matrix element $\langle A \rangle_{ij}$ is the number of seats on flight i required for itinerary j and, therefore, the column vector A^j specifies the total inventory required from the alliance network to satisfy itinerary j . In discussions below, we will assume that each request is for an individual passenger (i.e., $\langle A \rangle_{ij} \in \{0, 1\}$); however, group (multiseat) requests could be handled by creating additional columns with each positive element equaling the number of passengers in the group. For example, an itinerary from Rochester, NY to Denver, CO that passes through Chicago, IL would have ones in the rows for Rochester–Chicago and for Chicago–Denver. To handle a family of four looking to make the same trip, A would need another column with fours in those same rows.

For clarity, A can be partitioned as follows:

$$A = \begin{matrix} & \begin{matrix} N_1 & N_2 & N_S \end{matrix} \\ \begin{matrix} 1 \\ \vdots \\ m_1+1 \\ \vdots \\ m \end{matrix} & \left[\begin{array}{ccc|ccc|ccc} 1 & \cdots & n_1 & n_1+1 & \cdots & n_1+n_2 & n_1+n_2+1 & \cdots & n \\ \hline & A1 & & & & 0 & & & AS1 \\ \hline & & & & & & & & \\ \hline & & & & & A2 & & & AS2 \\ \hline \end{array} \right] \end{matrix},$$

such that the first n_1 columns have only positive elements in the first m_1 rows (airline 1's network), the next n_2 columns have only positive elements in the last m_2 rows (airline 2's network) and the final n_S columns have positive elements in both sets of rows.

Although interline itineraries may be sold by either alliance partner (requests for itineraries in N_S may be received by either airline 1 or 2), we assume that intra-line itineraries are only sold by the airline that operates the flights (requests for itineraries in N_c are only received by airline c). In practice, airlines do sell tickets for itineraries that are exclusively on another airline's network. With some additional notation, this possibility can be incorporated into the model and all

of the following results will continue to hold. To keep the exposition and notation simple, we will assume that each airline handles its own intraline requests.

The number of remaining (unsold) seats for flight i is denoted x^i . The m -dimensional vector \vec{x} is the joint vector of remaining inventory for the alliance: $\vec{x} \equiv \{x^1, \dots, x^{m_1}, x^{m_1+1}, \dots, x^m\}$.

3.1. The Demand Process

We consider a K -period booking horizon with the current period, denoted k , decreasing from K to 0. The probability that airline c receives a request for itinerary j in period k is $q_k^{cj} \geq 0$. We assume that each period is short enough such that the probability that the alliance receives more than one itinerary request in a given period is negligible. The probability that no request arrives is then

$$q_k^0 = 1 - \sum_{c \in \{1, 2\}} \sum_{j \in N} q_k^{cj} \geq 0. \quad (1)$$

The revenue R_k^{cj} associated with a request to airline c for itinerary j in a given period k , conditional on a request being made, is a nonnegative random variable with known cumulative distribution function (CDF) $F_k^{cj}(r)$. We assume that R_k^{cj} has a finite expectation. The complementary CDF is $\bar{F}_k^{cj}(r) = 1 - F_k^{cj}(r)$. Note that " c " in R_k^{cj} is the carrier that receives the consumer's request (the marketing airline). We assume that $F_k^{cj}(r)$ is differentiable with known density function $f_k^{cj}(r)$. However, wherever we express our results in terms of $f_k^{cj}(r)$, similar results can be found for non-continuous distributions.

3.2. Assumptions About the Arrival Process and Information Sharing

We assume that the distribution of each R_k^{cj} is independent of the realized revenue in preceding periods. Even simple (first-order) dependency, while theoretically easy to handle with our model, would be notationally and computationally cumbersome.

In addition, assume that each player formulates an open-loop dynamic program that does not use the realized arrival revenue stream as feedback for its optimization problem. One could imagine several closed-loop variations of our model. For example, demand intensity for a given itinerary could be characterized by an unknown parameter, which would be updated as demand is realized. Such models would be quite complex and, in practice, would likely be handled by updating the inputs to the model over the horizon without explicitly accounting for the future effects of this updating process when calculating the current value functions.

We also assume that there is independence between acceptance decisions in one period and the

arrival process in subsequent periods. Specifically, we assume that a customer, when denied a ticket, will not submit a new request to the alliance for the same or a similar itinerary. This assumption is consistent with the assumptions that underlie many of the models in the revenue management literature. Incorporating multiple customer preferences into the optimization problem is an area of ongoing research (e.g., see Talluri and van Ryzin 2004a). Within alliances, this behavior would add an interesting wrinkle to our problem because the revenue management decisions of each airline could affect the arrival process of its partners.

As noted in §1, in our model, the airlines share full information about their partners' inventory levels, forecasts of arrival processes, and revenue distributions. This allows the airlines to calculate, in each time period, a common expected value for a seat on any flight in the network. Using the terminology of game theory, we assume that each airline knows the strategies and payoffs of its partner and therefore plays a game of complete information. Although this model is stylized, it allows us to generate fundamental insights into the advantages and disadvantages of static and dynamic transfer price schemes.

Although we assume that each airline knows the *potential* payoffs of its partners, we do not assume that each airline immediately observes realized payoffs. Specifically, under the partner price scheme of §5.4.3, the operating airline must post its transfer prices for interline inventory without knowing the realized revenue associated with an interline request in that period (of course, the marketing airline sees any realized revenue). Therefore, the partners are playing a game of imperfect information. This assumption reflects an important source of information asymmetry found in the real world. For a given itinerary, there exists numerous classes and distribution channels through which the ticket can be sold; the range of prices across these classes and channels creates the distribution $F_k^{cj}(r)$ of revenue for each itinerary. Although the operating airline may know the distribution of revenue because prices are publicly posted, it cannot know the specific class being sold or channel being used at the moment the marketing airline receives a specific purchase request.

3.3. Assumptions About Revenue Sharing

In general, the proration scheme used by the alliance will influence both the total revenue received by the alliance and the allocation of revenues to each of the partners. We assume that the ultimate goal of each partner is to maximize its own wealth (revenue). It is reasonable to assume, however, that by forming an alliance, the partners are seeking a scheme that increases their joint profits, using some form of

ex ante revenue distribution (e.g., a participation fee) to ensure that all members of the alliance will continue to participate. In practice, this problem is often solved by finding a set of interline routes on each airline that leads to a rough balance in the revenue exchanged between the airlines (Ito and Lee 2007). The choice of mechanism for the distribution of total revenues is a bargaining problem that we do not examine here. We assume that some mechanism has already been chosen and that both airlines are willing to participate in the alliance. Therefore, our primary focus will be on examining how the various trading schemes affect total alliance revenue.

4. Centralized Control

Here, we describe the optimal policy for a single, centralized controller making all decisions to maximize total alliance revenue. In general, members of an airline alliance cannot adopt centralized revenue management controls (see the end of this section for further discussion), but these results are useful as they lead to an upper bound on the total revenue for the alliance. We will call this upper bound the *first-best* revenue.

4.1. Decision Process

The fundamental decision made by the centralized controller is whether to accept or reject a request for an itinerary j given the revenue offer R_k^{cj} and the current state of the system characterized by the remaining periods k and the remaining inventory of the alliance \vec{x} . Let $J_k(\vec{x})$ denote the total (current and future) expected value for the alliance given inventory \vec{x} with k remaining periods, and let $\Delta J_k(\vec{x}, A^j)$ be the opportunity cost to the alliance of the inventory required for itinerary j :

$$\Delta J_k(\vec{x}, A^j) \equiv J_k(\vec{x}) - J_k(\vec{x} - A^j). \quad (2)$$

For convenience, let $J_k(\vec{x}) = -\infty$ whenever one of the components of \vec{x} is negative. ($\vec{x}, \vec{0}, \vec{p}, \vec{u}$)

A policy for centralized control consists of a set of acceptance rules $u_k^j(\cdot; \vec{x})$ such that

$$u_k^j(r; \vec{x}) = \begin{cases} 1 & \text{if, at time } k \text{ with remaining inventory } \\ & \vec{x}, \text{ the alliance is willing to sell a} \\ & \text{ticket for itinerary } j \text{ with revenue } r, \\ 0 & \text{otherwise.} \end{cases}$$

We now define the joint arrival probability q_k^j and the corresponding conditional CDF F_k^j of the conditional revenue $R_k^j(r)$ for a request made to the *alliance*

(rather than to a particular partner c) for itinerary j in period k :

$$q_k^j = q_k^{1j} + q_k^{2j}$$

$$F_k^j(r) = \left(\frac{q_k^{1j}}{q_k^j}\right) F_k^{1j}(r) + \left(\frac{q_k^{2j}}{q_k^j}\right) F_k^{2j}(r).$$

The Bellman (1957) equations for optimal centralized control can then be written as

$$J_k(\vec{x}) = q_k^0 J_{k-1}(\vec{x}) + \sum_{j \in N} q_k^j E[R_k^j u_k^j(R_k^j, \vec{x}) + J_{k-1}(\vec{x} - A^j u_k^j(R_k^j, \vec{x}))],$$

$$J_0(\vec{x}) = 0 \quad \forall \vec{x} \geq \vec{0},$$

where

$$u_k^j(r, \vec{x}) = \arg \max_{u \in \{0, 1\}} \{ru + J_{k-1}(\vec{x} - A^j u)\}.$$

Given a request, the centralized controller either accepts the request (receives the associated revenue and reduces the inventory level) or denies the request and moves to the next period with the same inventory.

4.2. Optimal Policies

The decision faced by the centralized controller is identical to the decision faced by a single airline that maximizes the revenue generated by the combined network of the alliance. We can, therefore, apply results derived for a single airline network.

PROPOSITION 1. *The optimal acceptance policy for centralized control is of the form*

$$u_k^j(r; \vec{x}) = \begin{cases} 1 & \text{if } r \geq \Delta J_{k-1}(\vec{x}; A^j), \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. See Talluri and van Ryzin (1998), Proposition 1.

Under the optimal policy, the alliance accepts any request with associated revenue greater than or equal to the alliance's opportunity cost of the inventory used on that itinerary. Simply put, it accepts a request if it is beneficial (in expectation) to do so.

Practical limitations, however, prevent most alliances between large partners from ceding control of their revenue management systems to a central controller and using an optimal policy such as the one described in Proposition 1. Barriers to coordination include technical incompatibilities among revenue management systems within an alliance, competitive considerations (alliance partners are often competitors on many routes and therefore do not want to merge revenue management systems),

and antitrust laws. There are examples, however, of centralized control in the airline industry. Regional airlines sometimes allow their national partners to collect all revenues and make all booking decisions, and revenue-sharing is accomplished with a fixed payment per flight to the regional partner (for example, similar arrangements are used in the Continental/ExpressJet and United/Skywest alliances; see Shumsky 2006). For the remainder of this paper, we compare this centralized policy with the policies followed by airlines when revenue management decisions are distributed among the partners in the alliance. That is, the following decentralized control schemes have been used, or are intended for use, among major airlines such as the primary members of the SkyTeam, Star, and Oneworld alliances.

5. Decentralized Control

In this section, we examine airline behavior when revenue management decisions are decentralized among alliance partners. We assume that each alliance partner is free to accept or reject a request for an interline seat, as is true under the *free sale* system that is commonly used by major airline alliances (Boyd 1998a). In our model, interline sales require the following steps. First, an airline (hereafter the *marketing airline*) receives a request for an interline itinerary. Next, a transfer price is set for the seats on flights operated by its partner (the *operating airline*) that are needed to complete the itinerary (there are a variety of methods for setting transfer prices and we will describe specific schemes in §§5.3 and 5.4). Next, the operating airline decides whether to make its seat available and then the marketing airline decides whether to sell the complete itinerary. Finally, if the itinerary is sold, the transfer price is paid to the operating airline.

In §5.1, we will describe our model for the alliance under decentralized control. In §5.2, we will show by counterexample that no transfer pricing scheme can guarantee optimality under such a system, and we gain insights into the pitfalls inherent in transfer pricing schemes by examining the equilibrium behavior of the alliance partners under a generic decentralized scheme. In §§5.3 and 5.4, we will describe the equilibrium behavior of the partners under specific transfer price schemes. In §5.3, we examine static proration in which revenue from all interline tickets is split according to a fixed proportion. This scheme is currently used within many alliances. In §5.4, we analyze three dynamic schemes, which are based on systems proposed by Vinod (2005) and on systems that are being considered in the industry. In §5.5, we discuss the benefits of allowing the operating airline to set the transfer price and therefore share any surplus revenue received by the alliance for an interline request. In

§5.6, we consider how revenue is allocated between partners under each scheme and, in §5.7, we discuss the computational challenges one faces when attempting to calculate the equilibrium behavior of partners in an alliance.

5.1. Decision Process

We model the set of dynamic decisions for both airlines as a finite-horizon Markov game (Heyman and Sobel 2004). Although at the highest level the formation of the alliance can be viewed as a cooperative game, the contractual revenue-sharing mechanism must be implemented within each airline's revenue management system. These revenue management systems are inherently noncooperative, optimizing each airline's revenue without taking into account each decision's impact on the partner. Therefore, we assume that, given the transfer-pricing rules of the alliance, the revenue management systems of the airlines are locked in a noncooperative game. The two alliance airlines are the players in the game and, in §3.2, we described the information available to each player. The players' possible actions are quite simple: whether to accept or reject an itinerary request. In addition, under the *partner price* scheme described in §5.4.3, the operating airline has one more action—setting the transfer price. Because we use a Markov game, immediate payments and transition probabilities in each state depend only on the action in that state.

Let $J_k^c(\vec{x})$ denote the total (current and future) expected value for airline c given inventory \vec{x} with k remaining periods and with $J_k(\vec{x})$ denoting, as before, the total value for the alliance so that $J_k(\vec{x}) = J_k^1(\vec{x}) + J_k^2(\vec{x})$. As in Equation (2), the opportunity cost of the inventory used by an itinerary is denoted with a ΔJ term, although here we are concerned with each airline's individual opportunity cost:

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

in addition to the opportunity cost of the alliance as a whole $\Delta J_k(\vec{x}, A^j)$. As with centralized control, we define $J_k^c(\vec{x}) = -\infty$ whenever a component of \vec{x} is negative.

A policy for airline c consists of a set of acceptance rules $u_k^c(\cdot; \vec{x})$ such that

$$u_k^c(r; \vec{x}) = \begin{cases} 1 & \text{if, at time } k \text{ with remaining inventory } \\ & \vec{x}, \text{ airline } c \text{ is willing to sell a ticket} \\ & \text{for itinerary } j \text{ with net revenue } r, \\ 0 & \text{otherwise.} \end{cases}$$

Under the partner price scheme, the policy also includes setting the internal transfer price p_k^{cj} for each subitinerary.

The transfer price $p_k^{cj}(\vec{x})$ is a real number associated with each airline c , itinerary j , inventory level \vec{x} , and period k . For certain schemes, $p_k^{cj}(\vec{x})$ is also a function of the revenue associated with the request. To simplify the notation, however, we will not include R_k^{cj} as an argument of p_k^{cj} . Airline c 's partner must pay $p_k^{cj}(\vec{x})$ to airline c to sell the interline itinerary j . Let $\vec{p}_k^c(\vec{x})$ be the n -vector of all transfer prices in period k .

Note that we allow transfer prices to vary across each and every itinerary even if the subitinerary used on the operating airline is the same across multiple itineraries. A specific alliance arrangement may not allow for this level of detail. In particular, the marketing airline may request a subitinerary from the operating airline without revealing the entire itinerary. Therefore, within each period, the alliance will use a single transfer price for each subitinerary on the operating airline regardless of the itinerary being sold by the marketing airline. Although we do not examine the precise effects of this assumption, one would expect that a reduction in the amount of shared information would reduce the overall value of the alliance under decentralized control.

Although the specific form of the Bellman (1959) equations in the decentralized alliance will depend on the transfer price scheme used, the general form can be written as

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

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

The first summation corresponds to airline c 's intraline itinerary requests. As with the centralized model, airline c must then decide whether to accept a request. The second summation corresponds to airline c 's interline itinerary requests. Again, it must choose to accept or deny the request; however, the revenue on which this decision will be made is the revenue associated with the request less the transfer price paid to the alliance partner. The remaining two summations correspond to interline and intraline requests to airline c 's partner and the final term corresponds to the

“no arrival” case. Though airline c receives no revenue in the cases corresponding to the final summation, the change in its partner’s inventory does affect its future expected value.

Note that in (5.1), the accept/reject control variables u represent actions taken by the marketing airline and the formulation does not explicitly allow the operating airline to reject a request even though, under free sale, this action is available to the operating airline. We will see that it will not be necessary to explicitly model the operating airline’s acceptance policy under any of the dynamic schemes described in §5.4 because under all three schemes, the transfer price is always sufficiently large such that the operating airline will choose to accept the sale. Under the static schemes of §5.3, the operating airline may choose to reject a sale. In §5.3, we will discuss a modification to (5.1).

5.2. Nonoptimality of Markovian Transfer Price Schemes

Before examining specific transfer price schemes in detail, we describe a simple counterexample to demonstrate that no Markovian transfer scheme can guarantee network optimality as long as the transfer scheme is based solely on sales of interline itineraries. By “Markovian,” we refer to schemes that are completely defined by the current state of the network and do not depend on past states. Non-Markovian schemes that depend on the particular sample path (the history of which airline sold each seat, for how much, and when) could achieve optimality in the following counterexample. The complexity of such schemes, however, would make them impossible to implement.

Consider two airlines, 1 and 2, each operating one flight. Each flight has one remaining seat. Table 1 shows the expected demand over a two-period horizon. In the second column, an itinerary (x, y) requires “ x ” seats on airline 1 and “ y ” seats on airline 2. In the second to last period (period 2), each airline is equally likely to receive a request for its intraline itinerary with associated revenue of \$250. In the final period, airline 1 receives a request for an interline itinerary for \$400 with probability one. Clearly, it would be best for the alliance if the airline receiving the intraline request were to turn it down, leaving its inventory for the interline itinerary.

Let p be the transfer price in the final period paid to airline 2 if there is sufficient inventory remaining.

Table 1 Data for Counterexample to Transfer Pricing Optimality

Period k	Itinerary A^i	Marketing airline	Revenue (\$)	Probability
2	(1, 0)	Airline 1	250	0.5
	(0, 1)	Airline 2	250	0.5
1	(1, 1)	Airline 1	400	1

Therefore, at the beginning of period 2, the opportunity costs of the intraline inventory for airline 1 and airline 2 are $(\$400 - p)$ and p , respectively. Because each intraline request can be fulfilled without any of its partner’s inventory, each airline would maximize its own value by accepting an intraline request if its revenue exceeds its opportunity cost of its inventory (see Theorem 1 below for a formal proof of this behavior). Thus, to prevent either airline from filling an intraline request, p must satisfy both $p > \$250$ and $\$400 - p > \250 , or $\$250 < p < \150 , which is a contradiction.

Note that network optimality could be guaranteed if payments are made for intraline itineraries as well as interline itineraries. For example, assume that the airlines set $p = \$100$. Then, in period 2, let airline 1 offer \$151 to airline 2 if airline 2 agrees *not* to sell the intraline ticket if a request for that ticket arrives. Given such a subsidy scheme, neither airline will accept an intraline request and the network is optimized. Such transfer payments for intraline tickets, however, are impossible to implement for a variety of technological, competitive, and legal reasons.

Although no realistic Markovian transfer price scheme is universally optimal, certain schemes have intuitive appeal. For example, some practitioners have suggested that a seat’s opportunity cost (sometimes called its bid price) would be a logical transfer price (Vinod 2005). Although we will analyze each transfer price scheme separately, there are some common results worth noting. These results hold for all the schemes (static and dynamic) analyzed here. The results will also provide us with more general insights into why any transfer price scheme can fail to achieve first-best.

THEOREM 1. *For the Markov game described in §5.1, there exists a unique, pure strategy Markov perfect equilibrium in which the marketing airline adopts the policy*

$$u_k^{ij}(r; \vec{x}) = \begin{cases} 1 & \text{if } r \geq \Delta J_{k-1}^c(\vec{x}, A^i) + p_k^{-ij}(\vec{x}), \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. See Appendix A.

Theorem 1 shows that the marketing airline will accept any request that provides it with net revenue that exceeds its opportunity cost of the inventory used in the itinerary. The net revenue is the revenue received from the external customer for the itinerary minus any transfer price paid to the operating airline.

The counterexample illustrates an adverse consequence of the result in Theorem 1. Because there is no transfer price paid for the sale of an intraline itinerary, each airline makes intraline decisions without considering the effects of those decisions on its partner’s revenue. Therefore, even if a centralized

controller were to make all *interline* acceptance decisions (removing decision rights on interline itineraries from the marketing and operating carriers), the choice of the revenue-sharing method for interline itineraries would still affect the purely intraline decisions of the partners. This point is emphasized in Corollary 1 and the subsequent discussion.

COROLLARY 1. *The equilibrium control for intraline requests is of the form*

$$u_k^{c_j}(r; \vec{x}) = \begin{cases} 1 & r \geq \Delta J_{k-1}^c(\vec{x}, A^j), \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. This is immediate from Theorem 1 and $p_k^{-c_j}(\vec{x}) = 0$ for $j \in N_c$. □

The critical revenue level for airline c 's intraline decision is its own opportunity cost of the inventory used for the itinerary, much like the optimal decision for a single airline. In this case, however, the optimal (centralized) decision for the alliance—shown in Proposition 1—is determined by the total opportunity cost of the itinerary of *both* partners. That is, the critical value should be

$$\Delta J_{k-1}(\vec{x}, A^j) = \Delta J_{k-1}^c(\vec{x}, A^j) + \Delta J_{k-1}^{-c}(\vec{x}, A^j)$$

We refer to $\Delta J_{k-1}^{-c}(\vec{x}, A^j)$, the effect of the change in one airline's (here, the marketing airline's) inventory on its partner's value, as the *second-order effect*. We refer to $\Delta J_{k-1}^c(\vec{x}, A^j)$, the effect on the airline's own value, as the *first-order effect*. Our intuition is that inventory has a positive value; however, a simple example demonstrates that second-order effects—which correspond to the partner's value of the inventory—can be either positive or negative.

Consider the alliance shown in Figure 1 in which airline 1 operates flights A and B while airline 2 only operates flight C. Airline 1 offers itineraries AB and B (both intraline), and airline 2 offers AC (an interline itinerary).

Looking at the second-order effect (that on airline 2) of the sale of airline 1's itineraries, we expect oppositely signed values. The sale of a B itinerary frees up (in expectation) the A inventory needed by airline 2 to fill an AC request so we expect a negative opportunity cost for airline 2. That is, airline 2 is better off if




Figure 1 Sample Alliance for Control Comparisons

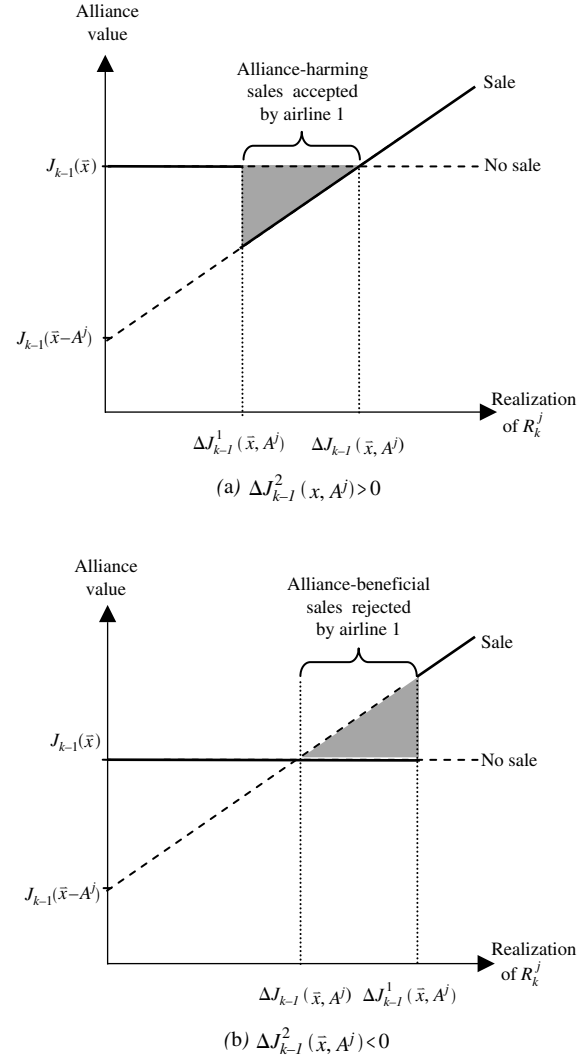


Figure 2 Loss of Potential Revenue from Inefficient Intraline Acceptance Policies

airline 1 sells a B itinerary. Conversely, a sale of an AB itinerary uses up A inventory so we expect a positive opportunity cost for airline 2; airline 2 is worse off if airline 1 sells an AB itinerary. Formally, we expect

$$\Delta J_{k-1}^2(\vec{x}, A^B) < 0 \quad \text{and} \quad \Delta J_{k-1}^2(\vec{x}, A^{AB}) > 0.$$

Therefore, airline 1 will overvalue its B itineraries, not selling them when it would benefit the alliance to do so, and will undervalue its AB itineraries, selling them when it does not benefit the alliance. Figure 2 illustrates the effect of these decisions on expected alliance revenue. The straight lines (horizontal and diagonal) are alliance values given that the itinerary request is accepted ("sale") or rejected ("no sale"). The bold lines indicate the acceptance policies in the decentralized alliance and alliance, losses are shown in gray.

Although no transfer price scheme can guarantee that the second-order effect will be incorporated in

each airline’s intraline decision making, we will show in §5.5 that certain transfer price schemes can indirectly reduce the impact of ignoring second-order effects, leading to more efficient intraline decisions.

5.3. Static Proration

In practice, revenue sharing for interline sales is often governed by static proration (SP) contracts that prorate the revenue received from an accepted request according to fixed proportions. (Such contracts are often enforced via relatively infrequent, ex-post sharing of revenue information so that the model formulated here is consistent with the information-sharing assumption described at the end of §3.2). One form of static proration specifies how revenue should be split for each and every itinerary. If airline c is the operating airline and carries a customer who paid the marketing airline r for itinerary j , airline c receives $\alpha^{cj}r$ as a transfer payment while the marketing airline retains $(1 - \alpha^{cj})r$ in revenue. To simplify the notation, we will assume that $\alpha^{-cj} = (1 - \alpha^{cj})$ so, in this case, the marketing airline r ’s share of the revenue is $\alpha^{-cj}r$ for itinerary j . We will refer to this form of static proration as *itinerary-specific SP*.

In practice, airlines sometimes use a common proration rate for multiple itineraries. The most extreme version uses a common (or *universal*) proration rate α for all itineraries. We examine two types of universal schemes, one based on the identity of the marketing airline and another that fixes the proportion for each airline and ignores whether an airline is the marketing or operating carrier. First, under “*universal SP (marketing)*,” the marketing airline receives $(1 - \alpha)r$ while the operating airline receives αr . In the following analysis, it will be useful to associate a proration rate with a particular airline c . Thus, under universal SP (marketing), $\alpha^c = \alpha$ if c is the operating airline and $\alpha^c = (1 - \alpha)$ if c is the marketing airline. Second, under “*universal SP (airline specific)*,” we assume that airline 1 receives αr and that airline 2 receives $(1 - \alpha)r$. Therefore, $\alpha^c = \alpha$ if $c = 1$ and $\alpha^c = (1 - \alpha)$ if $c = 2$. Table 2 summarizes the transfer prices paid to the operating airline under the static proration schemes. Table 2 also summarizes the dynamic transfer price schemes that will be described in detail in §5.4

Note that distinctions among the three static proration schemes will be relevant in the numerical examples described in §6. The results here, however, apply to all three static schemes. For convenience, we use the proration term α^{cj} throughout the following analysis, although the “ j ” can be eliminated for the universal schemes.

Table 2 Summary of Transfer Prices Paid to Operating Airline c

	Static proration (SP)	Dynamic transfer prices	
Itinerary-specific SP	$\alpha^{cj}R_k^{-cj}$	Bid price (full value)	$\Delta J_{k-1}^c(\bar{x}, A^j)$
Universal SP (marketing)	αR_k^{-cj}	Bid price (partial value)	$\Delta J_{k-1}(\bar{x}, A^{cj})$
Universal SP (airline specific)	αR_k^{-cj} to airline 1 and $(1 - \alpha)R_k^{-cj}$ to airline 2	Bid price proration	$\frac{\Delta J_{k-1}^c(\bar{x}, A^j)}{\Delta J_{k-1}(\bar{x}, A^j)} R_k^{-cj}$
		Partner price	Chosen by the operating airline (see Theorem 4)

THEOREM 2. For the marketing airline, the equilibrium interline acceptance policy under a static proration scheme is

$$u_k^{cj}(r; \bar{x}) = \begin{cases} 1 & \text{if } r \geq \Delta J_{k-1}^c(\bar{x}, A^j)/\alpha^{cj}, \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. See Appendix A.

The interline acceptance policy in Theorem 2 ensures that the marketing carrier earns at least its opportunity cost of the inventory used. Given static proration and the model formulation in (5.1), however, an interline itinerary may be accepted when its total revenue is less than the full alliance’s opportunity cost such that the operating airline earns less than the opportunity cost of its seats.

An obvious modification would be to give the operating airline as well as the marketing airline a veto over the sale of interline itineraries.

THEOREM 3. Under a static proration scheme, if both the marketing and operating airline may veto a sale, then in equilibrium an interline itinerary is sold only if

$$R_k^{cj} \geq \max\{\Delta J_{k-1}^c(\bar{x}, A^j)/\alpha^{cj}, \Delta J_{k-1}^{-c}(\bar{x}, A^j)/(1 - \alpha^{cj})\}.$$

PROOF. See Appendix A.

The acceptance criteria in Theorem 3 guarantees that no request is accepted that would hurt either partner, but increases the likelihood that some requests that are profitable for the alliance as a whole will be rejected. Under the widely used free sale system, both partners have the power to accept or reject a sale. Therefore, we focus on this modified model for the remainder of this section and for the numerical experiments in §6.

Now, define a *centrally optimal* decision as a decision in period k given inventory \bar{x} that would be optimal for the centralized system in that state. Note that centrally optimal interline decisions may not necessarily be optimal for the decentralized network because in the decentralized network, the intraline decisions may not be centrally optimal.

COROLLARY 2. *Under a static proration scheme, the alliance will, in equilibrium, make the centrally optimal interline acceptance decisions if*

$$\alpha^{cj} = \frac{\Delta J_{k-1}^c(\bar{x}, A^j)}{\Delta J_{k-1}(\bar{x}, A^j)}.$$

PROOF. After substituting this α^{cj} into the acceptance rule in Theorem 2 both terms in the maximization become $\Delta J_{k-1}(\bar{x}, A^j)$, the critical value in centralized control. \square

We term the ratio on the right-hand side of Corollary 2 as airline c 's *value ratio* for itinerary j in period k . Corollary 2 suggests that a static proration scheme will perform well if the proration rates are chosen such that

$$\alpha^{cj} \approx \Delta J_K^c(\bar{x}, A^j) / \Delta J_K(\bar{x}, A^j), \quad (4)$$

which is the value ratio at the start of the horizon, and if the ratio of opportunity costs is relatively stable for most demand realizations. Note that this ratio is actually a function of the proration rates α^{cj} , so a solution to (2) is a fixed point of the value functions J_K^c .

Static schemes are simple to implement, requiring relatively little information to be exchanged among partners after the initial proration-rate negotiation. In addition, by splitting the itinerary's revenue between the partners, such schemes do account for the changing values generated by entire itineraries. Theorem 3, however, identifies the flaw in any static scheme: The *relative* value of seats between the two partners may change over time and the static policy does not account for these changes. There are additional barriers to truly effective implementation of static policies. Finding the appropriate proration ratios for all individual itineraries can be a daunting task: Because the value functions J_K^c depend on the entire set of n_S proration rates, the system must be optimized over a continuous space with dimension $[0, 1]^{n_S}$, where n_S may be in the 1,000s. In addition, even if a small set of proration rates seems to work well across a particular network, static schemes have no internal mechanism to adjust for changing network parameters. For example, arrival rates and revenue distributions can change dramatically over time. After a significant change, the proration rates must be reoptimized and the alliance contract renegotiated. We will observe the impact of such changes in §6.

We now make one final point about the performance of static proration under a special case

COROLLARY 3. *When the alliance is composed entirely of interline itineraries, it can achieve first-best revenue by using static proration.*

PROOF. See Appendix A.

Note that Corollary 3 holds for all three forms of static proration if the proration rates are chosen correctly. For universal SP (airline specific), any proportion will work; for universal SP (marketing), the proportion must be 0.5; and for itinerary-specific SP, all itineraries must share the same proportion.

5.4. Dynamic Transfer Prices

In this section, we examine three revenue sharing schemes that dynamically change transfer prices to reflect the state of the system. These schemes are (i) *bid price*, where each airline simply charges its current opportunity cost or "bid price" of the seats in question and, therefore, the marketing airline retains all revenue in excess of its partner's current bid price; (ii) *bid price proration*, where the transfer price is adjusted so that the operating airline receives a share of the consumer revenue that is equal to its value ratio for the itinerary in question; and (iii) *partner price*, where each airline strategically chooses a price for its inventory that its partner must pay to use its subitinerary as part of an interline itinerary.

5.4.1. Bid Price Scheme. Under the *bid price* scheme, each airline posts its true opportunity cost (the bid price) for each interline itinerary in the alliance. If the marketing airline chooses to accept a request, then it must pay its partner the bid price as its transfer price, keeping any remaining revenue for itself. Conceptually, in the bid price scheme, the operating airline's value is unaffected by an interline sale because it is exactly compensated for the change in its inventory level. For now, we assume that the transfer price equals the operating airline's opportunity cost of *all* inventory sold in the itinerary. That is,

$$\begin{aligned} p_k^{cj}(\bar{x}) &= \Delta J_{k-1}^c(\bar{x}, A^j) \\ &= \Delta J_{k-1}^c(\bar{x}, A^{cj}) + \Delta J_{k-1}^c(\bar{x} - A^{cj}, A^{-cj}), \end{aligned} \quad (5)$$

where A^{cj} and A^{-cj} are the required subitineraries of itinerary j on airline c 's and $-c$'s networks, respectively.

The first term on the right-hand side of (5) is the first-order effect of the change in the operating airline's own inventory, and the second is the second-order effect of the marketing airline's change in inventory on the operating airline given that the first change has already been made. We will refer to this transfer price as a *full value* bid price.

THEOREM 4. *The equilibrium interline acceptance policy for a bid price scheme is centrally optimal if a full value bid price is used. The critical value in this case is*

$$\Delta J_{k-1}(\bar{x}, A^j).$$

PROOF. Add the critical acceptance policy and the transfer price, and apply Theorem 1 and Proposition 1. \square

Theorem 4 shows that the bid price scheme places no premium on the alliance’s opportunity cost, so its interline acceptance decisions are efficient (i.e., centrally optimal). However, we will show in §5.5 that use of a bid price scheme can lead to inefficient *intra-line* acceptance decisions, and these inefficient decisions may reduce total alliance revenue across the horizon.

There are a few implementation issues associated with the bid price scheme. The most important may be the problem of monitoring to ensure honest posting of bid prices. We will discuss this issue in §5.4.3. Another implementation issue is the choice of bid prices to use. In our analysis, we assumed that airlines agree to post their full opportunity costs for each itinerary. An alternative choice, previously suggested in Vinod (2005), would be to use only the operating airline’s opportunity cost of *its own* inventory in the itinerary i.e., neglecting the second term in (5.1). That is,

$$\begin{aligned} p_k^c(\vec{x}) &= \Delta J_{k-1}^c(\vec{x}, A^c) \\ &= \Delta J_{k-1}^c(\vec{x}, A^j) - \Delta J_{k-1}^c(\vec{x} - A^j, A^{-c}). \end{aligned}$$

This choice leads to inefficient acceptance decisions for interline itineraries because it ignores the second-order effects of the marketing airline’s change in inventory.

Another issue is how to handle the case when the operating airline’s opportunity cost is negative. Strict adherence to the policy would require the operating airline to post its true opportunity cost and thus subsidize the marketing airline’s sales. The operating airline would, therefore, pay the marketing airline to take its own inventory. It is unlikely that airlines would be willing to agree to this method; a logical alternative would be to simply post a zero price in place of any negative value. Doing so will produce inefficient interline acceptance decisions. In the numerical experiments of §6, we will use exact, full-value transfer prices for the bid price scheme and will allow negative transfer prices.

5.4.2. Bid Price Proration Scheme. In the *bid price proration* scheme, as in the bid price scheme, each airline posts its current opportunity cost for each interline itinerary that its partner sells. However, unlike the bid price scheme, if the marketing airline chooses to accept a request, then the revenue received is prorated by the value ratios of each airline. Specifically, the operating airline (airline c here) receives

$$\frac{\Delta J_{k-1}^c(\vec{x}, A^j)}{\Delta J_{k-1}^c(\vec{x}, A^j) + \Delta J_{k-1}^{-c}(\vec{x}, A^j)} r = \frac{\Delta J_{k-1}^c(\vec{x}, A^j)}{\Delta J_{k-1}(\vec{x}, A^j)} r, \quad (6)$$

and the marketing airline retains

$$1 - \frac{\Delta J_{k-1}^c(\vec{x}, A^j)}{\Delta J_{k-1}(\vec{x}, A^j)} r = \frac{\Delta J_{k-1}^{-c}(\vec{x}, A^j)}{\Delta J_{k-1}(\vec{x}, A^j)} r. \quad (7)$$

The marketing airline is free to fill the itinerary as long as the operating airline’s share of the revenue exceeds its posted opportunity cost.

THEOREM 5. *The equilibrium interline acceptance policy for the bid price proration scheme is centrally optimal with critical acceptance values $\Delta J_{k-1}(\vec{x}, A^j)$.*

PROOF. See Appendix A.

Theorem 5 shows that under the bid price proration scheme, the airlines respond to interline requests in a manner identical to their response under the bid price scheme—the interline requests are accepted if and only if it benefits the alliance to do so.

When the revenue associated with the request exactly equals the total opportunity cost of the partners—i.e., when there is no surplus revenue—then the bid price proration scheme gives each partner the same revenues as the bid price scheme. However, unlike the bid price scheme, each partner receives a share of the revenue proportional to its relative opportunity cost when the surplus revenue is greater than zero. The operating airline’s share can beneficially affect its decisions for intraline itineraries, resulting in higher revenues over the horizon than with the bid price scheme (see §5.5).

The bid price proration scheme also has implementation issues similar to those described for the bid price scheme. However, because proration rates are used to calculate transfer prices, the impact of zero and negative opportunity costs requires even more attention. A typical revenue-sharing arrangement would require these proportions to be between zero and one. Because opportunity costs can be negative, however, it is possible for the proportions to be negative, greater than one, or even infinite if the two opportunity costs add to zero. Therefore, it is impractical to implement the scheme exactly as described in Equations (6) and (7), and the partners must agree on methods to handle extreme cases. We propose two methods. The first method, like the bid price scheme, would replace negative opportunity costs with zero and “round” the proration rates to zero or one, ensuring that payments fall between zero and r . The second method would require an airline with a negative opportunity cost to subsidize its partner by exactly that amount instead of by an amount proportional to r , leading to payments in excess of r or payments in the “wrong direction” (i.e., the operating airline pays the marketing airline to use its seats). We use this second method in the numerical experiments of §6.

5.4.3. Partner Price Scheme. In the *partner price* scheme, each airline posts a dynamically updated list of transfer prices for each interline itinerary that its partner offers for sale. The partner price scheme can be used as a model of an actual contract that gives

the operating airline the power to set transfer prices. It is also important to see the partner price scheme as a model for a *bid price* scheme in which the partners “game” the system and are untruthful about their bid prices. We will see here that if bid prices cannot be monitored, each partner has an incentive to post higher prices than its actual bid prices. In our numerical experiments, we will see that such “gaming” of a bid price scheme can reduce alliance revenue. We will also see that for certain networks, the inflation of bid prices can actually increase alliance revenue.

Now, we discuss implementation details of our partner price scheme. Our model is consistent with the timing of two scenarios: Either all transfer prices are set at the beginning of each period, or the operating airline generates a transfer price on demand when the marketing airline makes a particular request. These transfer prices are based on the distribution of the revenues offered to the marketing airline because the operating airline does not know the actual realization of the revenue associated with a request when setting the prices (see the discussion at the end of §3.2). The marketing airline, on receiving a request, bases its acceptance decision on the realized value of the revenue and the posted transfer price. Note that the operating airline’s problem is to find the best transfer price for a “take it or leave it” offer, one of the “greed and regret” problems described by Sheopuri and Zemel (2007).

For Theorem 6, define $h_k^{cj}(r) = f_k^{cj}(r)/F_k^{cj}(r)$ and $g_k^{cj}(r) = rh_k^{cj}(r)$. Function g_k^{cj} is often called the generalized hazard rate function of F_k^{cj} .

THEOREM 6. *The equilibrium transfer price policy for each itinerary satisfies*

$$p_k^{cj} = \Delta J_{k-1}^c(\bar{x}, A^j) + 1/h_k^{-cj}(\Delta J_{k-1}^{-c}(\bar{x}, A^j) + p_k^{cj}).$$

If the distribution of R_k^{-cj} has increasing generalized failure rate (IGFR), i.e., $g_k^{-cj}(r)$ is nondecreasing, then the transfer price is guaranteed to be unique.

PROOF. See Appendix A.

Theorem 6 shows that the transfer price set by the operating airline (here, airline c) equals its opportunity cost for the itinerary plus a premium equal to the reciprocal of the hazard rate of the itinerary’s revenue evaluated at the threshold value of the marketing airline’s acceptance decision—the sum of the marketing airline’s opportunity cost and the transfer price. Uniqueness of the transfer price is then guaranteed if the revenue distributions have increasing generalized failure rate (IGFR), as is the case for the normal, uniform, exponential, beta, and many other distributions. The size of the premium chosen by the operating airline balances the increase in revenue received for each request accepted by its partner with the decreasing probability that the net revenue will be sufficiently high for its partner to accept.

COROLLARY 4. *The equilibrium acceptance policy under a partner price scheme has an acceptance threshold value:*

$$\Delta J_{k-1}(\bar{x}, A^j) + 1/h_k^{cj}(\Delta J_{k-1}^c(\bar{x}, A^j) + p_k^{-cj}(\bar{x})).$$

PROOF. Add the critical acceptance and transfer pricing policies and apply Theorem 1. □

Note that the strategic premium chosen by the operating airline to maximize its own expected revenue causes the marketing airline to make interline acceptance decisions that are *not* centrally optimal. The marketing airline will decline a request that would be beneficial to the alliance because the net revenue it receives does not fully compensate it for the opportunity cost of the inventory used. Figure 3 shows the expected revenue for the marketing (“M”) and operating (“O”) airlines with no premium versus a positive premium.

Although the premium leads to inefficient acceptance decisions for interline itineraries, we will show

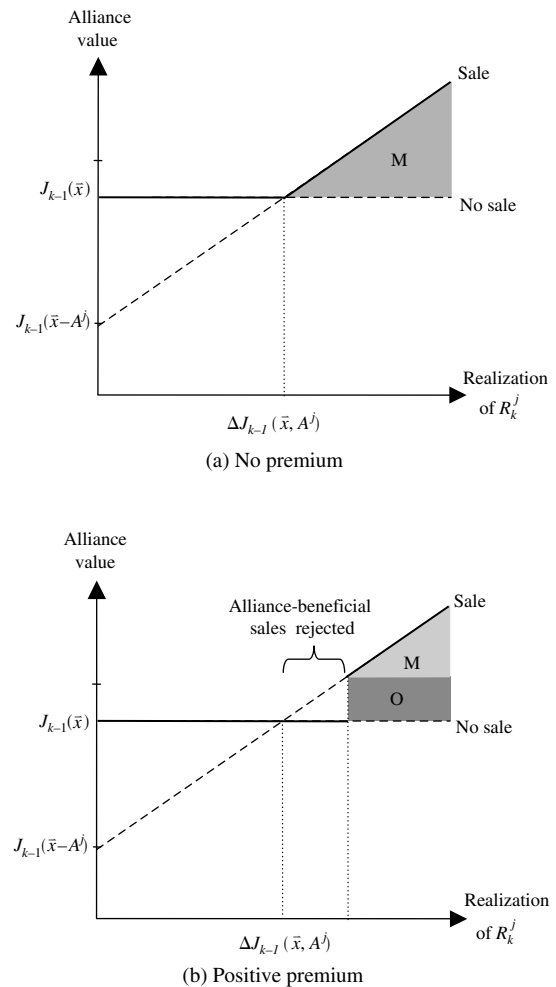


Figure 3 Alliance Value Under the Partner Price Scheme
 Note. The shaded areas are the expected surplus revenue received by the marketing (M) and operating (O) airlines.

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

in §5.5 that the premium may provide indirect benefits by improving *intra*line acceptance decisions.

5.5. Benefits of Surplus Sharing

Although the strategic premium in the partner price scheme can lead to inefficient interline decisions, the use of such a scheme may produce higher expected revenue over the horizon than other revenue-sharing alternatives. To see why, consider the example presented in Table 3.

In period one, airline 2 receives a request for its *intra*line itinerary with expected revenue of \$50 with probability one. If airline 2 has the required seat, it will always accept because its opportunity cost is \$0 in this period, giving an opportunity cost in period 2 of \$50. With the *bid price* scheme, airline 1 would accept its interline itinerary request for any revenue greater than \$50 (sum of the opportunity costs), paying airline 2 exactly its \$50 opportunity cost. In period 3, airline 2 would then accept any request with revenue greater than \$50—the opportunity cost of its *intra*line inventory in period 3—so it would sell its last seat, preventing the alliance from selling the valuable interline itinerary in period 2. The expected revenue for airlines 1 and 2 under this scheme is \$0 and \$100, respectively, with total expected revenue for the alliance of \$100.

Under the partner price scheme, the transfer price charged by airline 2 in period 2 will be \$775, leading to an acceptance level of \$775 for airline 1 and an opportunity cost of \$576 for airline 2 in period 3. As a result, airline 2 would not sell its *intra*line itinerary in period 3, leaving the inventory for sale by airline 1 in period 2. The expected revenues for airlines 1 and 2 under this scheme are \$273 and \$576, respectively, with total expected revenue for the alliance of \$839.

This example illustrates the benefits to the alliance of allowing its partner to share in the surplus revenue it receives for interline itineraries. Because airlines will ignore second-order effects when making *intra*line decisions, an airline that receives only its own opportunity cost for inventory under the *bid price* scheme will undervalue its inventory when, in the future, that inventory might be used by its partner to fill a valuable interline request. When operating airlines receive a portion of the expected surplus revenue, they consider the benefit of these future interline sales and (at least partially) account for alliance opportunity costs when making *intra*line decisions.

Table 3 Data for Example Illustrating Benefit of Premiums

Period k	Itinerary A^j	Marketing airline	Distribution of $R_k^{c^j}$	Probability $q_k^{c^j}$
3	(0, 1)	Airline 2	$U_{[50, 150]}$	1
2	(1, 1)	Airline 1	$U_{[500, 1,500]}$	1
1	(0, 1)	Airline 2	$U_{[25, 75]}$	1

The choice between partner price and bid price is, therefore, a tradeoff between the cost of inefficient interline decisions and the benefit of surplus sharing, which helps to convey information that leads to better *intra*line decisions. The bid price proration scheme maintains the efficient interline decisions of bid price while providing some of the information sharing of partner price. The relative costs and benefits of each scheme depend greatly on the parameters of the alliance, as illustrated in §6.

5.6. Revenue Allocation

Here, we discuss briefly the implications of each dynamic transfer price scheme on the allocation of revenue between alliance partners. The relative level of revenue received by the operating airline for an interline request will be higher under a partner price scheme than under a bid price scheme. In the partner price scheme, the operating airline is given more decision rights—it chooses the premium that benefits it the most. Therefore, it extracts a portion of the surplus revenue that the marketing airline can retain completely under the bid price scheme. The bid price proration scheme also provides the operating airline with a portion of the surplus revenue. Because the revenue is split according to the value of the inventory provided by each partner, it might be argued that this is a “fair” distribution. However, bid price proration sharing, like partner price sharing, does not reward the marketing airline for generating interline requests.

5.7. Computational Limitations

Given that the revenue management of a single airline network of realistic size requires the adoption of approximation methods (see Talluri and van Ryzin 2004b), computational limitations place tight restrictions on the size of the alliance problem that can be analyzed using the full dynamic program described earlier. For example, the size of the dynamic program’s state space for an alliance comprised of two equally sized airlines would be roughly the square of the size of the separate single-airline problems. Therefore, an important step from theory to practice will be the development of approximation methods for solving the alliance problem. The interaction of the partners within the game framework may preclude the direct use of existing approximation methods, or the application of existing methods may require significant transformations of the alliance problem described earlier.

6. Numerical Examples

To illustrate several of the effects discussed earlier, we examine the impact of each revenue-sharing scheme

on a set of sample alliances. In §6.1, we examine the relative performance of the three dynamic transfer-pricing schemes analyzed in §5 for two small alliances—one symmetric and one asymmetric—under changing system parameters. Then, in §6.2, we compare the performance of the best dynamic scheme, the partner price scheme (which can be interpreted as a dynamic bid price scheme with untruthful partners), and the three static schemes. As discussed in §5.7, computational requirements limit the size of the problems that we could analyze. Though the small sizes of these examples do not allow us to make definitive statements about the performance of these schemes in real-world networks, the examples do highlight the fundamental points discussed above.

6.1. Relative Performance of Dynamic Schemes

In this section, we examine two small alliances in which each airline operates one flight—A for airline 1 and B for airline 2—with 10 seats each. The time horizon has 30 periods. Revenues are normally distributed and increase stochastically as the time to departure approaches. See Appendix B for detailed information on the parameters.

The first example that we examine is a balanced (symmetric) alliance. Each partner sees the same demand distribution for its intraline itinerary (itinerary A or B) and an equal share of the interline requests (AB). The revenue distributions associated with the intraline and the interline requests are the same for both partners.

Figure 4 shows the results for this symmetric alliance as the proportion of intraline requests versus interline requests grows. The vertical axis represents the alliance revenue as a percent of the first-best revenue. At the far left, customers only request interline itineraries so the two schemes that make efficient interline decisions—the bid price proration scheme

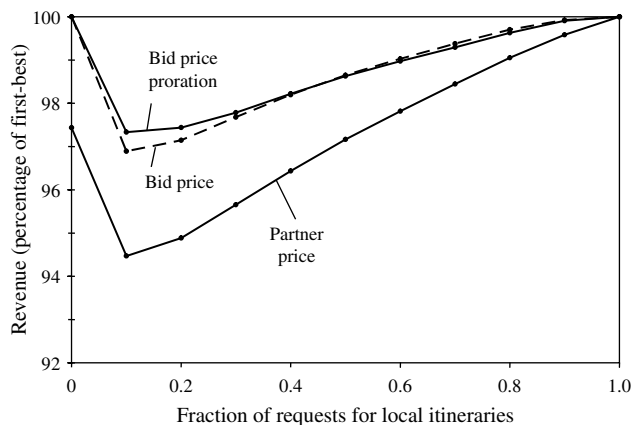


Figure 4 Percent of First-Best Total Revenue in Sample Symmetric Alliance

and the bid price scheme—perform optimally. The inefficiencies of the partner price scheme or, alternatively, the effects of cheating under the bid price scheme are also clear. Because each partner behaves in its own self-interest, transfer prices are too high and alliance revenues suffer. As in the prisoner's dilemma, this produces lower revenues for each partner because total alliance revenue is split 50/50 in this symmetric alliance.

On the far right, the alliance is receiving only intraline requests and, therefore, behaves as two separate airlines making individually optimal decisions. Because no interline requests are received, the choice of sharing scheme does not affect the total value and all schemes perform optimally. Finally, in the region between these two extremes, we see that the inefficiencies in both intraline decision making and the handling of interline requests reduce the performance of each scheme.

To further illustrate the benefits of surplus sharing discussed in §5.5, we now consider an asymmetric alliance in which only airline 1 receives requests for its intraline itinerary (A) while both airlines may receive requests for the interline itinerary (AB). Figure 5 shows the effects on alliance value as the fraction of interline requests received by airline 1 moves from zero to one. Figure 6 shows the percentage of alliance revenue extracted by airline 1 for each scenario shown in Figure 5.

At the far left of Figures 5 and 6, airline 1 receives all intraline requests while airline 2 receives all interline requests. We see in Figure 5 that the partner price scheme has the best performance while bid price proration outperforms the bid price scheme. On the left side of Figure 6, we see that under the partner price scheme, airline 1 receives a relatively large proportion of the alliance value by choosing a high transfer price. Bid price proration balances the contributions of both airlines while the bid price scheme awards more of

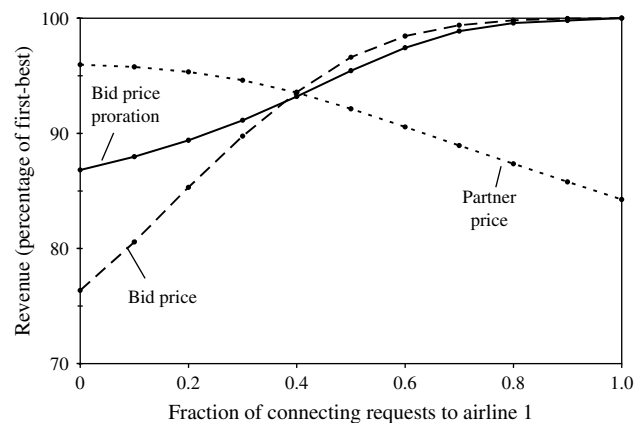


Figure 5 Percent of First-Best Total Revenue in Sample Asymmetric Alliance

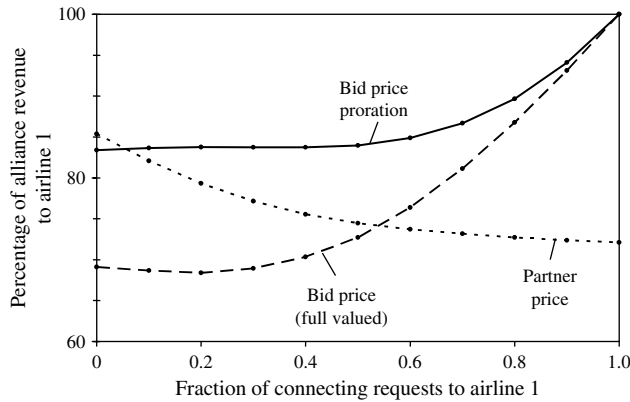


Figure 6 Percent of Alliance Value Captured by Airline 1 in the Asymmetric Alliance

the value generated by the interline itineraries to the marketing airline, airline 2.

The left sides of Figures 5 and 6 reinforce the insight of §5.5: Allowing the operating airline to set its own transfer price and capture more revenue can improve the performance of the alliance as a whole because this facilitates the sharing of information about the value of inventory. The example also shows that bid price proration can facilitate such information sharing but to a lesser degree. We should emphasize, however, that the left side of Figures 5 and 6 represents an extreme case in which all sales of the high-valued interline itineraries are made by a single airline, as might be the case in an alliance between a large international airline and a smaller national airline. When the airlines are closer to being equal partners, as in Figure 4, the partner price scheme reduces alliance revenue as well as the revenue of each partner.

Now, on the far right side of Figure 5, airline 2 does not bring any sales to the alliance so that under the (optimal) bid price proration and bid price schemes, airline 2’s opportunity cost is always zero. Therefore, airline 1 pays nothing per seat to airline 2 and captures 100% of the revenue (see Figure 6). The right side of Figure 5 suggests the benefits of the capacity purchase agreements that are typically used for national and regional alliances, in which the national and airline controls all revenue management activities, collects all revenue, and only pays a fixed price (say, a fee per flight) to the intraline alliance partner for its participation (see Shumsky 2006 for more details on these agreements).

6.2. Comparing Static and Dynamic Schemes

Here, we describe experiments that compare the performances of static and dynamic schemes in a series of networks with increasing heterogeneity in the value ratios across itineraries.

In the sample network for these experiments, airlines 1 and 2 both operate two flights: A and B for

airline 1, and C and D for airline 2. All flights have 4 seats remaining and the time horizon has 20 periods. Flights A and C travel into a connecting city while B and D fly out of that city. Therefore, there are two connecting itineraries: AD and CB. For our baseline model, the arrival probabilities and revenue distributions are chosen such that the interline value ratios for airline 1 are approximately 0.58 for both itineraries. That is, airline 1’s opportunity cost represents 58% of the total opportunity cost of the alliance for each of the interline itineraries. (To be more precise, the value ratios depend on the proration rate used. The value 0.58 and other value ratios described next are generated under the itinerary-specific SP scheme using optimal itinerary-specific proration rates. We provide more details on this scheme next.)

We induce heterogeneity in the value ratios by increasing (stochastically) the revenues and arrival probabilities of each airline’s intraline inbound itinerary with corresponding decreases in each airline’s intraline outbound itinerary. Specifically, the intraline itinerary revenue distributions and arrival probabilities are

$$\begin{aligned} E[R_k^{1A}] &= 250 + 50\beta & q_k^{1A} &= 0.1 + (0.1)\beta, \\ E[R_k^{1B}] &= 250 - 50\beta & q_k^{1B} &= 0.1 - (0.1)\beta, \\ E[R_k^{2C}] &= 150 + 50\beta & q_k^{2C} &= 0.1 + (0.1)\beta, \\ E[R_k^{2D}] &= 150 - 50\beta & q_k^{2D} &= 0.1 - (0.1)\beta. \end{aligned}$$

Figure 7 shows our results when the value ratio heterogeneity factor β is increased from zero to one with all other parameters held constant. At the extreme right, with $\beta = 1$, the value ratios of airline 1 have diverged from 0.58 up to 0.92 for AD and down to 0.25 for CB.

The line labeled “best dynamic” in Figure 7 shows the highest revenue achieved by the three dynamic schemes (the bid price scheme was the winner for all points shown in Figure 7). Figure 7 also shows the

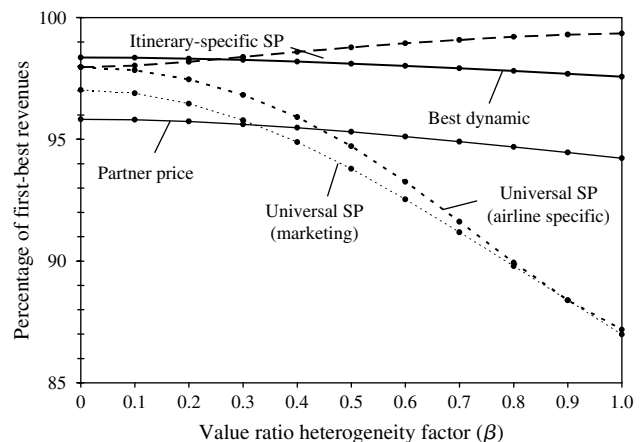


Figure 7 Percent of First-Best Under Increasing Value Ratio Heterogeneity

partner price scheme and allows us to see the impact if each partner distorts reported bid prices for its own (myopic) gain. The remaining lines show the revenues from the three static mechanisms described in §5.3: (1) a single proportion is used to determine airline 1's share of the revenue for all interline itineraries ("universal SP (airline specific)"); (2) a single proportion is used to determine the marketing airline's share of the revenue for all interline itineraries ("universal SP (marketing)"); and (3) a different proportion is used to determine airline 1's share of the revenue for each interline itinerary ("itinerary-specific SP"). For each static scheme, we numerically find the optimal proration ratios so Figure 7 reports the maximum revenue that can be achieved by each static scheme.

To understand the results in Figure 7, note that the value ratios determine the relative incentive needed by each airline to reserve its inventory for future interline requests rather than selling inventory for use in intraline itineraries. When an interline itinerary's value ratios are not 0.5, one airline values the inventory more than the other and, therefore, requires a larger share of revenues as incentive to make the right decision for the alliance as a whole. If the value ratios across the alliance are homogeneous (i.e., a given airline's opportunity cost represents the same proportion of the alliance opportunity cost for all itineraries), then a single ("universal") proration rate can be used effectively. Specifically, on the left side of Figure 7, a 58% share for airline 1 for both itineraries is optimal, although the performances of these optimal universal static schemes are slightly inferior to the performance of the best dynamic scheme. As the value ratios become heterogeneous, a universal proration rate cannot provide the proper incentives to maximize alliance revenues. On the right side of Figure 7, the performances of the universal static proration schemes fall substantially below the best dynamic scheme (bid price) and even well below the partner price scheme because the value ratios have diverged.

This observation applies to both universal static schemes. In Figure 7, the performance of the universal SP (marketing) scheme also demonstrates that though such a scheme could be used to provide an incentive to each airline to generate demand—an aspect of the problem not captured with our model—it fails to provide the proper incentives with regard to revenue management decisions.

As one would expect, the performance of the static scheme can be improved by choosing proration rates that are specific to each itinerary given their value ratios. Figure 8 shows the optimal itinerary-specific proration rates. In this example, using itinerary-specific rates allows static proration to perform significantly better under heterogeneous value ratios, even




Figure 8 Optimal Proration Rates for Itinerary-Specific SP Under Increasing Value Ratio Heterogeneity

outperforming the best dynamic scheme here. However, as discussed in §5.7, implementation of optimal itinerary-specific static schemes is difficult in large networks.

It should also be noted that the case used in Figure 7 does not include many elements found in practice that would be likely to hurt the performance of static schemes. To keep the number of parameters reasonable, in this example we used stationary revenue distributions; in practice, itinerary prices would be expected to change over the horizon. As discussed in §5.3, this second form of heterogeneity—over the horizon—would adversely affect the performance of static schemes. In a large number of additional experiments, we have found that the bid price and bid price proration schemes perform consistently better than static schemes. We now focus on the relative robustness of dynamic schemes.

Specifically, the performance of a static proration scheme declines as the characteristics of the alliance change, while proration rates are held constant. Figure 9 illustrates this point. To generate the figure, we first found the optimal proration rates, given $\beta = 0.50$, for both universal SP (airline-specific) and itinerary-specific SP. Then, we evaluated the performance of the static schemes using these proration rates in networks with other values of β . This scenario could represent a change in the system from an initial state ($\beta = 0.50$) that had been the baseline environment for an alliance agreement.

Because the best dynamic scheme automatically adjusts for these changes, its performance does not change substantially as β changes, ranging from 97.5% to 98.5% of first-best. Universal SP (airline-specific) falls far below the best dynamic scheme, dipping below 85% of first-best for $\beta = 1$. We also see that the itinerary-specific static proration




Figure 9 Percent of First-Best When Using Static Schemes That Are Optimal in a Network with $\beta = 0.5$

scheme, which is optimal when $\beta = 0.5$, performs poorly when value ratio heterogeneity is low ($\beta = 0$).

Another method to test the robustness of a static scheme is to use the wrong proration rate for a given network. In reality, there are many reasons why alliance members may choose an incorrect rate. For example, under commonly used mileage proration agreements, proration rates are based on the relative mileage of flight legs in an interline itinerary, an attribute that may not correspond with the actual relative value of inventory (see Boyd 1998a). In Figure 10, each curve shows the percent of first-best revenue achieved in a network with $\beta = 0.50$ under itinerary-specific SP, holding the proration rate for AD constant, and varying the rate for interline itinerary CB. The solid line shows results for the optimal AD proration rate of 77% (from the plot, we can see that the optimal pair of rates is CB = 39% and AD = 77%).

Although a small error (<10%) on either or both proration rates causes a reduction of less than 2% of first-best revenue, such an error can be sufficient to

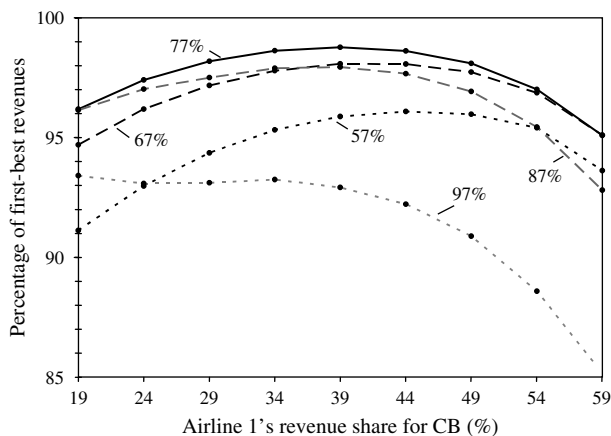


Figure 10 Percent of First-Best in a Network with $\beta = 0.5$ When Airline 1's Revenue Share for AD Varies from 57% to 97%

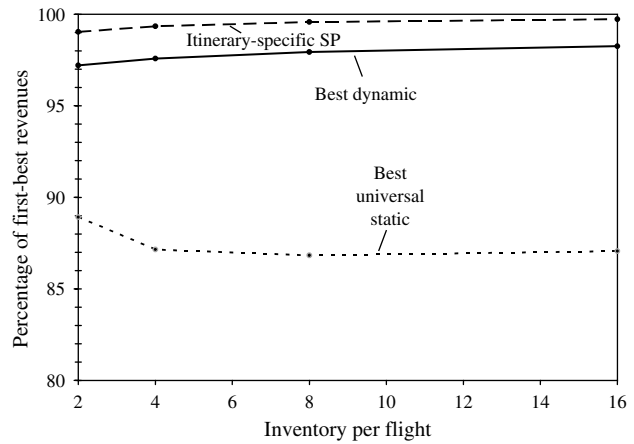


Figure 11 Percent of First-Best with Increasing Alliance Size

drop the performance of itinerary-specific SP below that of the best dynamic scheme. For example, at $\beta = 0.50$ in Figure 7, the best dynamic scheme achieves 98.11% of first-best revenue. At (87%, 49%) in Figure 10, itinerary-specific SP falls to 97.93% of first-best revenue. For more substantial deviations, itinerary-specific SP falls well below 90%.

Finally, we examined whether these results are artifacts caused by the small inventory levels on each flight. Figure 11 suggests that this is not the case. For these experiments, we use a network with $\beta = 1$. Inventory levels and the number of periods in the horizon are scaled from (2,10) up to (16,80), and Figure 11 shows the percentage of first-best alliance revenues for the universal SP (airline-specific), itinerary-specific SP and best dynamic schemes. The relative performances of the three schemes do not change significantly over this scale range, suggesting that similar results may hold for larger, more realistic inventory levels and horizon lengths.

7. Summary and Further Research

Airline alliances are selling increasing numbers of interline itineraries, creating the need to understand the impact of revenue-sharing mechanisms on alliance network performance. Although a few papers in the literature have addressed the broad challenges faced by airline alliances, this paper is the first to rigorously analyze the revenue-management behavior of alliance partners under various proration mechanisms. We began by defining a multiperiod, two-airline Markov game model of an alliance. Analysis of the model generated fundamental insights about the alliance revenue management problem. We find that no Markovian transfer-price scheme can guarantee optimal (first-best) revenues because the sharing of interline revenue cannot directly affect acceptance decisions for intraline itineraries. We then show that

the performance of static proration schemes depends on the homogeneity and stability of the relative values that each airline places on the inventory used in interline itineraries. Dynamic schemes adapt to changes in these “value ratios,” but there is no one best dynamic scheme. Dynamic schemes based on bid prices generate (centrally) optimal acceptance decisions for interline itineraries, but can fail to provide the operating airline with sufficient incentives to align intraline decisions. This alignment problem can be ameliorated by allowing the operating airline to set its own transfer prices (or, if the operating airline “games” a bid price system), but the resulting inflation of transfer prices can distort interline acceptance decisions.

We corroborated each of these insights using numerical experiments in small example networks. The experiments also served to reinforce observations about the advantages and challenges of implementing the static proration schemes that are currently used in the industry. In one example, we see that a static scheme with optimal itinerary-specific proration rates can slightly outperform the dynamic schemes considered here. Computing these optimal static proration rates and keeping them current, however, can be an enormous challenge for alliance partners, and our experiments show that the relative performance of the static schemes can decline significantly when the network parameters no longer match the proration rates.

These observations about computational complexity and practical considerations suggest two areas on which future research should focus: approximation methods and information asymmetry. As noted in §5.7, there has been significant work on approximating the single-airline network revenue management problem because of the problem’s size. In the single-airline problem, methods such as virtual nesting and additive bid prices are used to approximate the opportunity costs of an itinerary’s inventory. It would be natural to examine the impact of using similar approximation methods in an alliance.

In addition, in our model, we have assumed that the airline partners share demand and revenue forecasts as well as information about current inventory levels. As such, we assume that each can precisely calculate the value function for its partners, which is unlikely even for the smallest alliances. Future research should focus on the effects of relaxing this “complete information” assumption.

Another extension to our research would be to incorporate more advanced arrival and revenue processes and feedback into our model. As discussed in §3.2, dependency within the arrival stream could be added to the model, perhaps by reformulating the arrival process using a customer choice model. In addition, a closed-loop model could be created to update forecasts as demand is realized.

Finally, in our model, the airlines use a *free sale* scheme to exchange inventory; inventory is only exchanged when it is needed to fulfill a specific interline request. An alternate scheme is dynamic trading—called *soft blocking* by Boyd (1998a)—in which partners are allowed to “trade” (buy and sell) seats on each other’s flights throughout the horizon, shifting inventory to the partner that can use it to create the most value for the alliance. The potential benefits of dynamic trading are obvious, as demonstrated by Boyd (1998b), which describes conditions in which a static LP version of the problem can achieve the first-best revenue. Such a scheme, however, has many technical barriers to implementation. In addition, in a preliminary analysis, we have found that in a dynamic environment there are many conceptual and computational barriers that make airline behavior under dynamic trading difficult to predict. Describing the performance of dynamic trading will also be an interesting area of further research.

Acknowledgments

The authors thank Andy Boyd, Ben Vinod, Serguei Netesine, two anonymous reviewers, and an associate editor for their valuable input. The authors appreciate comments from participants in the Tuck School faculty seminar series and thank numerous analysts and revenue managers at Air France-KLM, Lufthansa, and United Airlines for providing information about the current and future practice of revenue management in alliances.

Appendix A.

Proofs of Markov perfect equilibrium and corresponding policies for decentralized alliance control schemes.

PROOF OF THEOREM 1 Given transfer prices $p_k^{-c}(\bar{x})$, each carrier’s Bellman (1957) equation is given by (3). The optimal control for the marketing airline is defined by

$$u_k^{c,j}(r, \bar{x}) = \arg \max_{u \in \{0, 1\}} \{(r - p_k^{-c}(\bar{x}))u + J_{k-1}^c(\bar{x} - A^j u)\}.$$

With the finite expectation assumption on the revenue distributions, this control can be rewritten as

$$u_k^{c,j}(r, \bar{x}) = \arg \max_{u \in \{0, 1\}} \{J_{k-1}^c(\bar{x}) + (r - p_k^{-c}(\bar{x}) - \Delta J_{k-1}^c(\bar{x}, A^j))u\}.$$

This is maximized by the function

$$u_k^{c,j}(r; \bar{x}) = \begin{cases} 1 & r \geq \Delta J_{k-1}^c(\bar{x}, A^j) + p_k^{-c}(\bar{x}), \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

which is the acceptance policy described in Theorem 1.

Because only one revenue request can arrive within each period, both $\Delta J_{k-1}^c(\bar{x}, A^j)$ and $p_k^{-c}(\bar{x})$ are independent of $\bar{u}_k^{-c}(\bar{x})$, which are airline $-c$ ’s acceptance decisions (as a marketing airline) in period k . Therefore, the best-response acceptance decisions for airline c , $\bar{u}_k^c(\bar{x})$, are constant with respect to the partner’s decisions as a marketing airline. That is,

$$u_k^{c,j}(r; \bar{x}, \bar{u}_k^{-c}) = u_k^{c,j}(r; \bar{x}).$$

Because the distribution of R_k^{cj} is continuous, (1) defines a unique best response by each airline; there are no “ties,” where $r = \Delta J_{k-1}^c(\bar{x}, A^j) + p_k^{-cj}(\bar{x})$.

If the transfer prices are sufficiently large in all cases so that the operating airline will always choose to accept a sale, then the marketing airline’s policies $u_k^{cj}(r; \bar{x})$ define a unique equilibrium. If the operating airline $-c$ may reject a sale, then let $\hat{u}_k^{-cj}(r; \bar{x})$ be the operating airline’s acceptance policy. An argument identical to the one above shows that the optimal policy for the operating airline is

$$\hat{u}_k^{-cj}(r; \bar{x}) = \begin{cases} 1 & p_k^{-cj}(\bar{x}) \geq \Delta J_{k-1}^c(\bar{x}, A^j), \\ 0 & \text{otherwise.} \end{cases}$$

Again, for each state of the system, this policy is unique and is independent of the partner’s action. Therefore, the pairs of best responses $[u_k^{cj}(r; \bar{x}), \hat{u}_k^{-cj}(r; \bar{x})]$ define a unique equilibrium. This unique equilibrium exists for each subgame defined by each time period and inventory state. Therefore, these best responses define a unique, pure strategy Markov perfect equilibrium. \square

Although throughout the paper we assume that the distribution of R_k^{cj} is continuous, if we allow a probability mass to exist in the revenue distribution, then one could imagine a mixed-strategy equilibrium. Suppose that the realized revenue $r = \Delta J_{k-1}^c(\bar{x}, A^j) + p_k^{-cj}(\bar{x})$ has a nonzero probability and consider the following optimal mixed strategy: Requests are accepted by the marketing airline with some probability γ and rejected with probability $1 - \gamma$ when that particular revenue is realized. The value to airline c , however, would be unaffected by any choice of γ because the expected value is the same in both cases. The independence of the accept or reject decisions would, again, ensure that the partners’ decisions would be unaffected as well. As such, there would be no obvious benefit to such a policy and the added complexity is likely to be prohibitive. Therefore, even if the distribution of R_k^{cj} is not continuous, one would expect that the airlines would play pure strategies.

PROOF OF THEOREM 2 (STATIC PRORATION). For this scheme, for a given share proportion α^c for airline c , the conditional Bellman (1957) equations take the form

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

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

where $u_k^{cj}(r, \bar{x}) = \arg \max_{u \in \{0,1\}} \{ru + J_{k-1}^c(\bar{x} - A^j u)\}$. Therefore, when $j \in N_s$, $u_k^{cj}(\alpha^{cj} R_k^{cj}, \bar{x}) = 1$ if and only if $\alpha^{cj} R_k^{cj} \geq \Delta J_{k-1}^c(\bar{x}, A^j)$. Because there is at most one arrival in a period, each

policy $u_k^{cj}(r, \bar{x})$ is independent of the decisions for other interline itineraries in the same period. Therefore, the threshold strategy described above yields a Nash equilibrium and the acceptance rules form a Markov perfect equilibrium. \square

PROOF OF THEOREM 3 (MODIFIED STATIC PRORATION). Under this scheme, both airlines may veto an interline acceptance decision. If both partners accept an interline itinerary, the marketing airline $-c$ receives its share of the revenue $(1 - \alpha^{cj})R_k^{-cj}$ while the operating airline receives $\alpha^{cj}R_k^{-cj}$. By Theorem 2, the operating airline will accept the request if and only if $\alpha^{cj}R_k^{-cj} \geq \Delta J_{k-1}^c(\bar{x}, A^j)$, and the marketing airline will accept the request if and only if $(1 - \alpha^{cj})R_k^{-cj} \geq \Delta J_{k-1}^c(\bar{x}, A^j)$. Therefore, the combined acceptance rule is

$$R_k^{-cj} \geq \max\{\Delta J_{k-1}^c(\bar{x}, A^j)/\alpha^{cj}, \Delta J_{k-1}^c(\bar{x}, A^j)/(1 - \alpha^{cj})\}. \quad \square$$

PROOF OF COROLLARY 3 (STATIC PRORATION). Proof by induction. First, let $H_k(\bar{x})$ denote the first-best revenues for the alliance in period k with remaining inventory \bar{x} . In period 0, any alliance has a value of zero. For an arbitrary period $k - 1$, $k \geq 1$, make the following assumption:

$$J_{k-1}^c(\bar{x}) = \alpha^c H_{k-1}(\bar{x}) \quad c \in \{1, 2\}.$$

It follows then that

$$\Delta J_{k-1}^c(\bar{x}) = \alpha^c \Delta H_{k-1}(\bar{x}) \quad c \in \{1, 2\}.$$

Any interline itinerary will be accepted if

$$\begin{aligned} R_k^c & \geq \max\{\Delta J_{k-1}^c(\bar{x}, A^j)/\alpha^c, \Delta J_{k-1}^c(\bar{x}, A^j)/(1 - \alpha^c)\} \\ & = \max\{\alpha^c \Delta H_{k-1}(\bar{x}, A^j)/\alpha^c, (1 - \alpha^c) \Delta H_{k-1}(\bar{x}, A^j)/(1 - \alpha^c)\} \\ & = \Delta H_{k-1}(\bar{x}, A^j). \end{aligned}$$

Thus, the alliance makes the same acceptance decision as the centralized controller. Additionally, because the share of revenues received by airline c is α^c , our inductive assumption holds with

$$J_k^c(\bar{x}) = \alpha^c H_k(\bar{x}) \quad c \in \{1, 2\}. \quad \square$$

PROOF OF THEOREM 5 (BID PRICE PRORATION). Assume that airline c has received a request for interline itinerary j in period k with a fare of r . Under bid price proration, the transfer price that carrier c must pay if it decides to accept the itinerary request is then

$$p_k^{-cj} = \frac{\Delta J_{k-1}^c(\bar{x}, A^j)}{\Delta J_{k-1}^c(\bar{x}, A^j)} r.$$

Provided that this transfer price is used, airline c ’s optimal acceptance policy is as given in Theorem 1 and the airlines’ acceptance policies constitute a Markov perfect equilibrium. Note that with this transfer price, we have

$$\begin{aligned} r \geq \Delta J_{k-1}^c(\bar{x}, A^j) + p_k^{-cj} & \Leftrightarrow r \geq \Delta J_{k-1}^c(\bar{x}, A^j) \\ & \Leftrightarrow p_k^{-cj} \geq \Delta J_{k-1}^c(\bar{x}, A^j). \end{aligned}$$

In other words, the marketing carrier accepts the itinerary if and only if it is beneficial for the alliance as a whole, and this is if and only if the operating carrier receives at least its opportunity cost in compensation. This proves Theorem 5 and also shows that the marketing and operating carriers’ interests are aligned when it comes to making interline acceptance decisions. \square

Appendix B. Table of Parameters for Sample Alliances

Parameter	Symmetric alliance			Asymmetric alliance	
Number of periods	30			30	
Number of seats on leg operated by airline 1	10			10	
Number of seats on leg operated by airline 2	10			10	
Probability of an arrival in each period	1			1	
Itinerary	A	AB	B	A	AB
Fraction of requests (for the symmetric alliance, “intra-line” is the fraction of requests for A or B, and we vary “intra-line” from 0 to 1)	0.5 * intra-line	1-intra-line	0.5 * intra-line	0.5	0.5
Share of requests to airline 1	100%	50%	0%	100%	Varies
Mean revenue per request (last period)	200	500	200	200	400
Growth in mean revenue over all periods (%)	20	20	20	20	20
Standard deviation of revenue per request	20	50	20	20	40

PROOF OF THEOREM 6 (PARTNER PRICE). Recall that we assume that the operating carrier knows only the distribution of the revenue associated with the request received by the marketing carrier, not the actual revenue offered by the prospective passenger. The operating carrier also knows both its own and the marketing carrier’s opportunity cost for the itinerary. In this scheme, the operating carrier chooses the transfer prices for an interline itinerary request received by the marketing carrier.

For the remainder of the proof, because the operating airline is making the decision of interest, we will refer to it as c and the marketing airline as $-c$. Thus, the relevant transfer prices established by the operating airline are $p_k^{cj}(\bar{x})$. Given these transfer prices, the marketing carrier’s Bellman (1957) equation is given by (3) and the optimal control is described in Theorem 1. Now, we consider the equilibrium transfer price chosen by the operating carrier. Assume that marketing airline $-c$ received a request for interline itinerary $j \in N_s$ and the operating carrier c has chosen a transfer price of p . Then, the conditional (on such a request having been received) expected total current and future revenues for airline c are given by

$$J_{k-1}^{cj}(\bar{x}) + \int_0^\infty (p - \Delta J_{k-1}^c(\bar{x}, A^j)) u_k^{-cj}(r - p, \bar{x}) dF_k^{-cj}(r) = J_{k-1}^{cj}(\bar{x}) + (p - \Delta J_{k-1}^c(\bar{x}, A^j)) \tilde{F}_k^{-cj}(p + \Delta J_{k-1}^c(\bar{x}, A^j)).$$

Taking the derivative with respect to p gives

$$\begin{aligned} \tilde{F}_k^{-cj}(p + \Delta J_{k-1}^c(\bar{x}, A^j)) - (p - \Delta J_{k-1}^c(\bar{x}, A^j)) f_k^{-cj} & \cdot (p + \Delta J_{k-1}^c(\bar{x}, A^j)) \\ = \tilde{F}_k^{-cj}(p + \Delta J_{k-1}^c(\bar{x}, A^j)) [1 - (p - \Delta J_{k-1}^c(\bar{x}, A^j)) h_k^{-cj} & \cdot (p + \Delta J_{k-1}^c(\bar{x}, A^j))], \end{aligned}$$

where $h_k^{-cj}(r) = f_k^{-cj}(r) / F_k^{-cj}(r)$. Hence, the optimal transfer price p_k^{cj} satisfies

$$p_k^{cj} = \Delta J_{k-1}^c(\bar{x}, A^j) + 1/h_k^{-cj}(\Delta J_{k-1}^c(\bar{x}, A^j) + p_k^{cj}).$$

Note that the condition of Theorem 3 is equivalent to

$$g_k^{-cj}(\Delta J_{k-1}^c(\bar{x}, A^j) + p_k^{cj}) = \frac{p_k^{cj} + \Delta J_{k-1}^c(\bar{x}, A^j)}{p_k^{cj} - \Delta J_{k-1}^c(\bar{x}, A^j)},$$

where $g_k^{-cj}(r) = r h_k^{-cj}(r)$ is the generalized hazard rate function of the itinerary’s revenue distribution $F_k^{-cj}(r)$. The RHS of this last

equation decreases from $+\infty$ when $p_k^{cj} = \Delta J_{k-1}^c(\bar{x}, A^j)$ to one when $p_k^{cj} \rightarrow \infty$. Now, assume that the distribution of R_k^{-cj} has increasing generalized failure rate (IGFR), i.e., $g_k^{-cj}(r)$ is nondecreasing. Given that $g_k^{-cj}(r)$ has a finite expectation, it follows from Theorem 2 of Lariviere (2006) that $\lim_{r \rightarrow \infty} g_k^{-cj}(r) > 1$. Hence, the operating airline’s optimal transfer price is unique.

Finally, because there is at most one arrival per period, the optimal pricing and acceptance policies for a particular interline itinerary j are independent of other policies for other itineraries in the same period. Therefore, the results in this section of Appendix A imply that under the IGFR assumption, the transfer price policies in Theorem 3 and acceptance policies in Theorem 1 form a Markov perfect equilibrium. \square

References

Agarwal, R., O. Ergun. 2007. Network design and allocation mechanisms for carrier alliances in liner shipping. Working paper. Georgia Institute of Technology, Atlanta.

Barron, R. II. 1997. Code share agreements: A developing trend in U.S. bilateral aviation negotiations. *Indiana Law J.* 72 529–543.

Belden, T. 2007. Southwest Airlines wants to add international service. *Philadelphia Inquirer*. May 11, City Ed., D03.

Bellman, R. 1957. *Dynamic Programming*. Princeton University Press, Princeton, NJ.

Belobaba, P. 1989. Application of probabilistic decision model to airline seat inventory control. *Oper. Res.* 37(2) 182–197.

Boyd, E. A. 1998a. Airline alliances. *OR/MS Today* (October) 28–31.

Boyd, E. A. 1998b. Airline alliance revenue management. PROS Strategic Solutions technical report, PROS Revenue Management, Houston.

Brueckner, J. 2001. The economics of international codesharing: An analysis of airline alliances. *Internat. J. Indust. Organ.* 19(10) 1475–1498.

Brueckner, J., W. T. Whalen. 2000. The price effects of international airline alliances. *J. Law Econom.* 43 503–545.

Cachon, G. P., S. Netessine. 2004. Game theory in supply chain analysis. D. Simchi-Levi, S. D. Wu, Z.-S. Shen, eds. *Handbook of Quantitative Supply Chain Analysis: Modeling in the eBusiness Era*. Kluwer, New York.

Cachon, G. P., P. H. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Sci.* 45(7) 936–953.

Fernandez de la Torre, P. E. 1999. Airline alliances: The airline perspective. Master’s thesis, Massachusetts Institute of Technology, Cambridge.

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at <http://journals.informs.org/>.

- Glover, F., R. Glover, J. Lorenzo, C. McMillan. 1982. The passenger mix problem in the scheduled airlines. *Interfaces* 12(3) 73–79.
- Granot, D., S. Yin. 2008. Competition and cooperation in decentralized push and pull assembly systems. *Management Sci.* 54(4) 733–747.
- Heyman, D. P., M. J. Sobel. 2004. Stochastic models in operations research. *Stochastic Optimization*, Vol. II. Dover Publications, Mineola, NY.
- Houghtalen, L., O. Ergun, J. Sokol. 2007. Designing mechanisms for the management of carrier alliances. Working paper, Georgia Institute of Technology, Atlanta.
- International Air Transport Association (IATA). 2007. Declared interline prorate values. Working paper, http://www.iata.org/nr/rdonlyres/f51dc9a0-4f23-4ae1-b586-617ec8ca6ded/0/dipv_final_conceptv1_0.pdf.
- Ito, H., D. Lee. 2007. Domestic code sharing, alliances, and airfares in the U.S. airline industry. *J. Law Econom.* 50(2) 355–380.
- Kasilingam, R. 1996. Air cargo revenue management: Characteristics and complexities. *Eur. J. Oper. Res.* 96 36–44.
- Lariviere, M. A. 2006. A note on probability distributions with increasing generalized failure rates. *Oper. Res.* 54(3) 602–604.
- Lariviere, M. A., E. L. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing Service Oper. Management* 3(4) 293–305.
- Littlewood, K. 1972. Forecasting and control of passenger bookings. *AGIFORS Sympos. Proc.* 12 96–117.
- Lott, S. 2006. Alliances expected to grow even larger. *Aviation Week Space Tech.* 164(19) 44.
- McGill, J., G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* 33(2) 233–256.
- Nagarajan, M., Y. Bassok. 2008. A bargaining framework in supply chains: The assembly problem. *Management Sci.* 54(8) 1482–1496.
- Nagarajan, M., G. Sošić. 2008. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *Eur. J. Oper. Management* 187(3) 719–745.
- Netessine, S., R. Shumsky. 2004. Revenue management games: Horizontal and vertical competition. *Management Sci.* 51(5) 813–831.
- Park, J. 1997. The effect of airline alliances on markets and economic welfare. *Transportation Res. Part E* 33(3) 181–195.
- Sheopuri, A., E. Zemel. 2007. The greed and regret (GR) problem. Working paper OM-2007-1, Stern School of Business, New York University, New York.
- Shumsky, R. A. 2006. The Southwest effect, airline alliances, and revenue management. *J. Revenue Pricing Management* 5(1) 83–89.
- Spengler, J. 1950. Vertical integration and antitrust policy. *J. Political Econom.* 58(4) 347–352.
- Talluri, K., G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Sci.* 44(11) 1577–1593.
- Talluri, K., G. van Ryzin. 2004a. Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* 50(1) 15–33.
- Talluri, K., G. van Ryzin. 2004b. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- U.S. Department of Transportation. 2006. Airline origin and destination survey (DB1B). Research and Innovative Technology Administration, Bureau of Transportation Statistics, Accessed February 13, 2006, <http://www.transtats.bts.gov/DataIndex.asp>.
- Vinod, B. 2005. Alliance revenue management. *J. Revenue Pricing Management (Practice Papers)* 4(1) 66–82.
- Vulcano, G., G. van Ryzin, C. Maglaras. 2002. Optimal dynamic auctions for revenue management. *Management Sci.* 48(11) 1388–1407.
- Wynne, R. 1995. From Queen Airs to RJs: Evolution in major-regional airline alliances and code sharing. *Handbook of Airline Economics*. Aviation Week Group, McGraw-Hill, New York.
- You, P. S. 1999. Dynamic pricing in airline seat management for flights with multiple legs. *Transportation Sci.* 33(2) 192–206.