

The Price Effects of International Airline Alliances

by

Jan K. Brueckner
Department of Economics
and
Institute of Government and Public Affairs
University of Illinois at Urbana-Champaign

W. Tom Whalen
Department of Economics
University of Illinois at Urbana-Champaign

December 1998

Abstract

This paper provides evidence of the effect of international airline alliances on fares. The main finding is that alliance partners charge interline fares that are 18–28 percent below those charged by nonallied carriers. According to our theoretical model, the main source of this fare reduction is internalization of a negative externality that arises from uncoordinated choice of interline “subfares” in the absence of an alliance. The paper also looks for evidence of an anticompetitive alliance effect in the gateway-to-gateway markets. While the point estimates show that an alliance between two previously-competitive carriers would raise fares by 4-6 percent, this effect is not statistically significant.

The Price Effects of International Airline Alliances

by

Jan K. Brueckner and W. Tom Whalen*

1. Introduction

Deregulation of U.S. airlines in 1978 spurred the adoption of a number of major innovations by the industry. Route structures were reorganized into hub-and-spoke networks, frequent flier programs were instituted, and computerized reservations systems were built. The latest innovation sweeping the industry is the formation of international airline alliances.

The oldest of the major alliances is operated by Northwest and KLM, and it was followed by an alliance between United and Lufthansa, which has since expanded to embrace 4 additional carriers, including Air Canada (it is known as the Star alliance). Delta also has major alliances with Swiss Air, Sabena, and Aeromexico, while American Airlines is allied with Canadian Air and is also pursuing a link with British Airways, which has yet to receive regulatory approval.

Through international alliances, U.S. carriers are able to extend their networks overseas without operating additional flights. The carriers can thus overcome restrictions on international service codified in various bilateral agreements, which still govern much overseas traffic despite a number of new “open skies” agreements. While circumventing such limitations, alliances also allow U.S. carriers to extend their global reach without investing new resources.

Since U.S. airlines, in conjunction with foreign carriers, could always provide “interline” service to any overseas destination, the need for alliances may not be obvious. Their appeal to the airlines arises from several sources. First, by coordinating flight schedules and ensuring gate proximity at connecting airports, the alliance partners can offer greater convenience to the passenger. Alliance travel thus resembles single-airline service, avoiding many of the inconveniences of a traditional interline trip. This effect, which attracts passengers away from nonallied carriers, is reinforced by unification of the partners’ frequent flier programs. In addition, the single-airline nature of alliance travel is often formalized in a codesharing agreement, where the component flights are ticketed as if they occurred entirely on one partner airline.

The second advantage of an alliance lies in the realm of pricing. Since the major alliances enjoy antitrust immunity, the alliance partners can engage in cooperative pricing of interline trips. By contrast, pricing of traditional interline travel is best viewed as the result of non-cooperative behavior, as argued in the theory developed below. Because of this difference, the theory predicts that the interline fares charged by alliance partners are lower than those of nonallied carriers. Thus, the theory says that, in addition to offering greater passenger convenience than a traditional interline trip, alliance travel is more attractively priced.

In some cases, alliances may generate anticompetitive effects that tend to offset the above benefits. These effects are illustrated by the controversy surrounding the proposed American/British Airways (AA/BA) alliance. The problem is that AA and BA currently compete on nonstop routes between U.S. “gateway” cities, such as New York and Chicago, and the major U.K. gateway at London-Heathrow. By jointly serving these “gateway-to-gateway” routes, the two carriers control 60 percent of U.S.-U.K. traffic. If the AA/BA alliance were approved and granted antitrust immunity, the carriers would be free to collude in setting prices in these nonstop markets. Government regulators fear that the resulting loss of competition would push up fares, to the detriment of U.S.-U.K. passengers. Such losses would tend to offset the benefits accruing to alliance passengers who make interline trips, passing through the gateway airports without stopping. These passengers do not suffer a loss of competition from the alliance, but instead enjoy the benefits of cooperation in pricing and scheduling, as noted above.

The purpose of the present paper is to investigate these issues by presenting empirical evidence on the price effects of international airline alliances. Using fare data from the U.S. Department of Transportation, we document the impact of alliances on interline fares, testing the above hypothesis of a negative effect. In addition, we explore the fare impact of overlapping service by alliance partners in gateway-to-gateway markets, looking for evidence of the fare escalation that has preoccupied regulators. These exercises are carried out in Sections 3 and 4 of the paper. Using the estimates, we present a simulation analysis of the proposed AA/BA alliance in subsection 4.5.

To motivate the empirical work, Section 2 presents a theoretical model, built on the previous work of Brueckner (1997) and Brueckner and Spiller (1991), that illustrates the favorable

effect of alliances on interline fares. The model focuses on the impact of cooperative pricing, abstracting from any convenience gains due to alliances, which may affect demand. The key assumption of the model, which follows Brueckner (1997), is that nonallied carriers set interline fares through noncooperative choice of “subfares.” The subfare gives the amount earned by a carrier for its portion of an interline trip, with the sum of the subfares giving the interline fare. Such unrestricted fare-setting behavior is consistent with the discussion of Doganis (1985), who indicates that previous formula-based methods of fare determination are now little used. The key implication of the model is that the carriers’ noncooperative subfare choices generate negative externalities. Alliances internalize these externalities through a revenue-sharing arrangement, and the result is a lower fare.

Another key feature of the model is the operation of hub-and-spoke networks by the carriers, with the goal of exploiting economies of traffic density. These economies, whose existence is documented by Caves, Christensen and Tretheway (1984) and Brueckner and Spiller (1994), arise because high traffic densities allow carriers to operate larger, more efficient aircraft and to disperse fixed costs over more passengers. In the presence of economies of density, international traffic gains from cooperative pricing of interline trips lead to a lower marginal cost of carrying an additional passenger within the networks of the alliance partners. This cost reduction puts further downward pressure on the alliance fare, as well as allowing the partners to charge lower fares for purely domestic travel.

In addition to the analysis of Brueckner (1997), the only previous theoretical work on alliances is presented by Park (1997). Empirical work on alliances is also sparse. Oum, Park and Zhang (1996) present an empirical study that shows how a codesharing agreement between two small carriers on an international route affects the market leader (data is from trans-Pacific routes). Park and Zhang (1998) offer an analysis of the fare impacts of alliances on North Atlantic routes, using published fare data rather than actual transaction data of the type used below. Finally, Youssef and Hansen (1994) present a case study of the discontinued SAS/Swissair alliance, showing its effects on service coordination and market concentration. This small recent literature on alliances extends a huge previous literature on the airline industry, which is surveyed by Morrison and Winston (1995).

2. A Model of Airline Alliances

2.1. An earlier model

Before developing the model, it is helpful to discuss the previous analysis of international alliances presented by Brueckner (1997). In his model, two airlines operate the route structure shown in Figure 1. Airline 1 operates the routes shown as dotted lines, using city H as its hub, while airline 2 operates the routes represented by solid lines, using city K as its hub. Both airlines serve the (transatlantic) city-pair market HK, which connects the hub (or gateway) cities. In addition to serving this gateway-to-gateway market, airline 1 serves three domestic city-pair markets, AH, BH, and AB, with service in the latter provided by connecting flights through the hub H (airline 2 provides parallel domestic service). Together, airlines 1 and 2 also provide interline service in the international markets AD, AE, BD, and BE, where passengers travel across the networks of both airlines in making their trip. These city-pair markets are denoted “behind-the-gateway” markets, reflecting the nonhub status of their endpoints. Finally, airline 1 provides “online” (single-carrier) service in the international markets AK and BK, again through its hub (airline 2 provides parallel EH and DH service).

In the absence of an alliance, airlines 1 and 2 compete in the gateway-to-gateway market, HK. They also determine the interline fare in the behind-the-gateway markets, again in noncooperative fashion. The carriers noncooperatively choose “subfares” for their portion of each interline trip, with the sum of the subfares determining the overall fare in the market. When the carriers operate instead as alliance partners, they collude rather than compete in the gateway-to-gateway market, which tends to raise the fare in that market. Collusion also occurs in behind-the-gateway markets, but instead of being anticompetitive, this behavior generates benefits. Since they split the revenue from such markets, the carriers cooperatively choose an overall fare instead of setting individual subfares. This cooperation internalizes the negative externalities from subfare choices, and the result is downward pressure on interline fares.

Thus, the expectation is that an alliance leads to higher fares in the gateway-to-gateway market and lower interline fares. While this outcome is not guaranteed in general, Brueckner’s numerical analysis of a parameterized version of the model shows that it emerges for almost all parameter values.

2.2. Setup of the current model

Although the above model is suggestive, it is not ideal for motivating an empirical study of airline alliances. The problem is that an international passenger making an interline trip can usually choose between several carrier pairs, either allied or nonallied. The model, however, assumes the absence of interline competition, with service provided only by a single carrier pair. To provide analysis that is empirically more relevant, the present model has four airlines rather than two. Two competing carriers then operate in each set of domestic markets, and the international markets are served by competing carrier pairs, either allied or nonallied.

To prevent this added complexity from making the analysis unmanageable, the model's route structure, shown in Figure 2, is simplified. In particular, the gateway-to-gateway market is suppressed by assuming that both carriers operate out of the same hub, H. This is literally impossible in an international context, but the resulting simplification allows empirical predictions regarding interline fares to be derived in a straightforward fashion. As can be seen in Figure 2, carriers 1 and 2 compete in the domestic city-pair markets AH, BH and AB, with carriers 3 and 4 competing in their domestic markets. In contrast to Figure 1, there are now just four international markets, AD, AE, BD, and BE, all of which require an interline trip. For simplicity, we assume that carriers 1 and 3 always form an interline pair, with 2 and 4 forming the other interline pair. The airlines within each of these competing pairs can operate as alliance partners or as nonallied carriers.¹

In contrast to Brueckner (1997) and many other models in the literature, we assume that passengers exhibit brand loyalty to particular carriers, implying that airlines sell differentiated products. Because of brand loyalty, a carrier (or carrier pair) can raise its fare above that charged by its competitor without losing all of its traffic. As will be seen, the analysis cannot proceed without this assumption. Formally, we assume that when carriers i and k both serve a city-pair market, charging fares of p^i and p^k , carrier i 's traffic is given by $q^i = D(p^i, p^k)$, with k 's traffic given by $q^k = D(p^k, p^i)$. $D(\cdot)$ is the symmetric demand function, which is decreasing in its first argument and increasing in the second. For example, $q_{AH}^1 = D(p_{AH}^1, p_{AH}^2)$ gives carrier 1's traffic in market AH when its fare is p_{AH}^1 and carrier 2's fare is p_{AH}^2 , with $p_{AH}^1 D(p_{AH}^1, p_{AH}^2)$ giving revenue. Note that the variable q_{AH}^1 measures traffic in both directions in market AH,

being equal to the number of passengers travelling from A to H and back, *plus* the number of passengers travelling from H to A and back, all on carrier 1. Finally, since we assume symmetry across city-pair markets in the demand for travel, the function $D(\cdot)$ applies to all domestic and international markets.

In the absence of an alliance, the interline fares are generated via the noncooperative choice of subfares, as in Brueckner (1997). In city-pair market AD, for example, airline 1's subfare is denoted s_{AD}^1 , while the subfare of its interline partner, carrier 3, is denoted s_{AD}^3 . The fare in the market is then $s_{AD}^1 + s_{AD}^3$, and carrier 1's AD traffic is $q_{AD}^{13} = D(s_{AD}^1 + s_{AD}^3, s_{AD}^2 + s_{AD}^4)$, where the notation q_{AD}^{13} indicates that this also represents carrier 3's traffic. Note that this traffic level depends on the subfares charged by the competing airline pair, which consists of carriers 2 and 4. Carrier 1's revenue in the market is $s_{AD}^1 D(s_{AD}^1 + s_{AD}^3, s_{AD}^2 + s_{AD}^4)$, while carrier 3's revenue equals D multiplied by s_{AD}^3 instead of s_{AD}^1 .

To write carrier 1's total revenue, the symmetry of city pairs can be exploited. Given symmetry, carrier 1's fares and traffic levels will be the same in markets AH and BH, allowing these variables to be written as p_{XH}^1 and q_{XH}^1 , where X refers to either A or B. Similarly, airline 1's subfares and traffic levels will be same in each of the international markets AD, AE, BD, BE, allowing these variables to be written as s_{XX}^1 and q_{XX}^{13} , with the latter traffic level applying also to carrier 3 (note that the first X subscript denotes A or B, while the second denotes D or E). Analogous notation applies to the other carriers. With these conventions, airline 1's revenue is $2p_{XH}^1 q_{XH}^1 + p_{AB}^1 q_{AB}^1 + 4s_{XX}^1 q_{XX}^{13}$, or

$$2p_{XH}^1 D(p_{XH}^1, p_{XH}^2) + p_{AB}^1 D(p_{AB}^1, p_{AB}^2) + 4s_{XX}^1 D(s_{XX}^1 + s_{XX}^3, s_{XX}^2 + s_{XX}^4), \quad (1)$$

where the multiplicative factors 2 and 4 serve to sum revenues across identical markets.

Now consider the case where each carrier pair (1 and 3, 2 and 4) operates as an alliance. The alliance partners set a total fare in each international market and split the market revenue. Fares are symmetric across international markets, and are denoted p_{XX}^{13} and p_{XX}^{24} for the two alliances. Carrier 1's revenue is now given by $2p_{XH}^1 q_{XH}^1 + p_{AB}^1 q_{AB}^1 + 2p_{XX}^{13} q_{XX}^{13}$, or

$$2p_{XH}^1 D(p_{XH}^1, p_{XH}^2) + p_{AB}^1 D(p_{AB}^1, p_{AB}^2) + 2p_{XX}^{13} D(p_{XX}^{13}, p_{XX}^{24}). \quad (2)$$

The last term in (2) comes from dividing the total international revenue earned by the alliance of carriers 1 and 3, equal to $4p_{XX}^{13}q_{XX}^{13}$, between the two partners.

Carrier 1's costs are equal to the cost of operating the two spokes of its network. The cost of operating a single spoke is $c(Q)$, where Q is the traffic density (i.e., total traffic) on the spoke. With economies of density, this cost function satisfies $c' > 0$ and $c'' < 0$, implying that cost per passenger is decreasing in traffic density. Since density is the same on both of carrier 1's spokes, being given by $Q = q_{XH}^1 + q_{AB}^1 + 2q_{XX}^{13}$, the cost of operating its network equals

$$2c(q_{XH}^1 + q_{AB}^1 + 2q_{XX}^{13}). \quad (3)$$

2.3. Equilibrium conditions

Airline 1's goal is to maximize profit, which equals the difference between revenue, given by (1) or (2), and the cost expression (3). In the non-alliance case, where revenue is given by (1), airline 1 makes decisions independently of all the other carriers. It chooses values for p_{XH}^1 , p_{AB}^1 , and s_{XX}^1 in Cournot fashion, taking the values of the other airlines' decision variables as given. The resulting first-order conditions are

$$q_{XH}^1 + p_{XH}^1 \frac{\partial q_{XH}^1}{\partial p_{XH}^1} - c' \frac{\partial q_{XH}^1}{\partial p_{XH}^1} = 0 \quad (4)$$

$$q_{AB}^1 + p_{AB}^1 \frac{\partial q_{AB}^1}{\partial p_{AB}^1} - 2c' \frac{\partial q_{AB}^1}{\partial p_{AB}^1} = 0 \quad (5)$$

$$2q_{XX}^{13} + 2s_{XX}^1 \frac{\partial q_{XX}^{13}}{\partial s_{XX}^1} - 2c' \frac{\partial q_{XX}^{13}}{\partial s_{XX}^1} = 0. \quad (6)$$

To interpret these conditions, consider (5). After dividing through by $\partial q_{AB}^1 / \partial p_{AB}^1$, the last term equals $2c'$, which gives the marginal cost of serving a passenger in market AB (the sum of marginal costs on the two network spokes used for the trip). After division, the first two terms in (5) represent the marginal revenue earned from an AB passenger (as a function of q_{AB}^1). The condition thus requires equality between marginal revenue and marginal cost. The other conditions are interpreted similarly.

Airline 3 solves a symmetric problem, and as a result, it chooses the same subfare as airline 1. Therefore, the total interline fare charged by carriers 1 and 3 is given by $p_{XX}^{13} = 2s_{XX}^1$. Condition (6) from above can then be rewritten as

$$2q_{XX}^{13} + p_{XX}^{13} \frac{\partial q_{XX}^{13}}{\partial p_{XX}^{13}} - 2c' \frac{\partial q_{XX}^{13}}{\partial p_{XX}^{13}} = 0, \quad (7)$$

where the fact that $\partial q_{XX}^{13} / \partial s_{XX}^1 = \partial q_{XX}^{13} / \partial p_{XX}^{13}$ is used.

When carriers 1 and 3 are allied, they confer in the choice of the interline fare p_{XX}^{13} . Since the carriers solve symmetric choice problems, they agree on the level of the fare, and as a result, only carrier 1's problem need be considered. The first-order conditions for the problem are given by (4), (5) and

$$q_{XX}^{13} + p_{XX}^{13} \frac{\partial q_{XX}^{13}}{\partial p_{XX}^{13}} - 2c' \frac{\partial q_{XX}^{13}}{\partial p_{XX}^{13}} = 0. \quad (8)$$

Equations (7) and (8) differ only in the coefficient of q_{XX}^{13} , which is smaller in the alliance case, where (8) applies. Since marginal revenue (as a function of p_{XX}^{13}) is therefore lower when the carriers are allied, we would expect the first-order condition to be fulfilled at a smaller value of p_{XX}^{13} than in the non-alliance case. In other words, we would expect airlines 1 and 3 to charge a lower fare as alliance partners than as nonallied carriers. To see the intuitive reason for this outcome, observe that in the absence of an alliance, an increase in carrier 1's subfare depresses carrier 3's profit by reducing traffic in the given market. While carrier 1 has no incentive to take this negative externality into account, revenue sharing under an alliance causes the externality to be internalized, moderating carrier 1's pursuit of a higher fare.

To establish this conclusion formally, the first step would be to recognize that because of symmetry, the choices of competing carriers (and carrier pairs) will be identical, an outcome that emerges in the both the alliance and nonalliance cases. Removing the carrier superscripts, the endogenous variables then become p_{XH} , p_{AB} , and p_{XX} . The equilibrium values of these variables in the nonalliance case are determined by (4), (5), and (7), while their equilibrium values in the alliance case are determined by (4), (5), and (8). These systems differ only in the

q_{XX}^{13} coefficient in (7) and (8), denoted ϕ , which has the value 2 in the nonalliance case and 1 in the alliance case. The effect of alliance formation can then be derived by comparative-static analysis with respect to the parameter ϕ .

Unfortunately, it is not possible to sign the resulting comparative-static derivatives without additional assumptions. It can be shown, for example, that if economies of density are absent, with marginal cost equal to a constant δ , then $\partial p_{XH}/\partial\phi = \partial p_{AB}/\partial\phi = 0$ and $\partial p_{XX}/\partial\phi > 0$.² Thus, the international fare is lower in the alliance case, as expected, while domestic fares are unaffected (recall that the alliance case has the lower ϕ).

To generate determinate results in the presence of economies of traffic density, it is helpful to impose specific functional forms for cost and demand. Such assumptions also allow analysis of the asymmetric case, where one carrier pair operates as an alliance while the other does not, a case that is of great empirical relevance (see section 4.3 below). Following Brueckner and Spiller (1991), the first assumption imposes linearity on the marginal cost function, which is given by $c'(Q) = 1 - \theta Q$, where $\theta > 0$ (the cost function itself is then quadratic). The second assumption is that brand loyalty to carriers (or carrier pairs) takes a particular form. In a representative market where carriers i and k compete, we assume that a given passenger will fly on carrier i if the fares satisfy $p^i < p^k + a$, or $a > p^i - p^k$, where a gives the individual's dollar-denominated preference for carrier i . If a is uniformly distributed over the interval $[-\alpha/2, \alpha/2]$, then airline 1's traffic is equal to

$$q^i = \int_{p^i - p^k}^{\alpha/2} \frac{1}{\alpha} da = \frac{1}{2} - \frac{p^i - p^k}{\alpha}. \quad (9)$$

Carrier k 's traffic is then

$$q^k = \frac{1}{2} + \frac{p^i - p^k}{\alpha}. \quad (10)$$

A restrictive feature of this formulation is that total demand is insensitive to the level of fares, always being equal to one. Fare differences serve only to divide this fixed total between carriers.

From (9) and (10), the demand derivatives in the above first-order conditions all equal $-1/\alpha$. Thus, a tight distribution of consumer preferences (a small α) makes demand highly

price sensitive. After substituting this derivative along with the marginal cost expression, the first-order conditions can be written

$$q_{XH}^1 - \frac{1}{\alpha} p_{XH}^1 + \frac{1}{\alpha} [1 - \theta(q_{XH}^1 + q_{AB}^1 + 2q_{XX}^{13})] = 0 \quad (11)$$

$$q_{AB}^1 - \frac{1}{\alpha} p_{AB}^1 + \frac{2}{\alpha} [1 - \theta(q_{XH}^1 + q_{AB}^1 + 2q_{XX}^{13})] = 0 \quad (12)$$

$$\phi q_{XX}^{13} - \frac{1}{\alpha} p_{XX}^{13} + \frac{2}{\alpha} [1 - \theta(q_{XH}^1 + q_{AB}^1 + 2q_{XX}^{13})] = 0, \quad (13)$$

where ϕ equals 1 in the alliance case and 2 in the nonalliance case. Since the equilibria are symmetric, as explained above, the competitors split the traffic in each domestic and international market. Therefore, all the traffic variables in (11)–(13) equal $\frac{1}{2}$, and fares can be computed directly. As in the situation where economies of density are absent, the equilibrium fares in the domestic markets XH and AB are the same with and without alliances. These domestic fares, which are the same across carriers, are given by

$$p_{XH}^* = \frac{\alpha}{2} + 1 - 2\theta \quad (14)$$

$$p_{AB}^* = \frac{\alpha}{2} + 2 - 4\theta. \quad (15)$$

Observe that the AB fare is higher than the XH fare, reflecting the higher marginal cost of a connecting trip.³ Note also that, by symmetry, (14) and (15) give the fares charged by carriers 3 and 4 in markets DH/EH and DE.

International fares in the alliance and nonalliance cases are given by

$$p_{XX}^{ally} = \frac{\alpha}{2} + 2 - 4\theta, \quad p_{XX}^{non-ally} = \alpha + 2 - 4\theta \quad (16)$$

As expected, (16) shows that formation of airline alliances leads to lower international fares. Also, observe that the alliance fare is the same as the AB fare in (15), reflecting the equivalence of the pricing problems in the two types of markets in the alliance case.⁴

2.4. The asymmetric case

The case that is perhaps most relevant to the empirical work is asymmetric, and it therefore differs from those considered so far. In this case, carriers 1 and 3 form an alliance, while carriers 2 and 4 provide interline service without forming an alliance. The fare charged by 1 and 3 is thus set jointly, while fare charged by 2 and 4 is the result of noncooperative choice of subfares. The equilibrium in this asymmetric case is determined as follows. The first set of equilibrium conditions applies to carrier 1, and it consists of (11)–(13) with $\phi = 1$. The second set of conditions applies to carrier 2, and it consists of (11)–(13) with the carrier superscripts changed from 1 to 2 and from 13 to 24. In addition, ϕ is set equal to 2 to reflect the absence of an alliance. The resulting six-equation system is then augmented with 6 demand relationships based on (9) and (10), yielding a 12-equation system that determines 6 fares and 6 traffic levels for carriers 1 and 2. The fare solutions are given by

$$p_{XH}^1 = \frac{12\alpha + 6\alpha^2 - 48\theta - 49\alpha\theta + 96\theta^2}{12(\alpha - 4\theta)} \quad (17)$$

$$p_{XH}^2 = p_{XH}^1 + \frac{\alpha\theta}{6(\alpha - 4\theta)} \quad (18)$$

$$p_{AB}^1 = \frac{12\alpha + 3\alpha^2 - 48\theta - 37\alpha\theta + 96\theta^2}{6(\alpha - 4\theta)} \quad (19)$$

$$p_{AB}^2 = p_{AB}^1 + \frac{\alpha\theta}{3(\alpha - 4\theta)} \quad (20)$$

$$p_{XX}^{13} = \frac{16\alpha + 5\alpha^2 - 64\theta - 54\alpha\theta + 128\theta^2}{8(\alpha - 4\theta)} \quad (21)$$

$$p_{XX}^{24} = p_{XX}^{13} + \frac{\alpha^2 - 2\alpha\theta}{8(\alpha - 4\theta)}. \quad (22)$$

By symmetry, the domestic fares charged by carriers 3 and 4 are given by (17)–(20).

To compare fares in the various markets, the sign of the common denominator expression $\alpha - 4\theta$ must be established. Using the second-order condition for the carrier maximization problems, it can be shown that this expression must be positive.⁵ Algebraic manipulation

then establishes the following conclusions:

$$p_{XH}^1 < p_{XH}^* < p_{XH}^2 \quad (23)$$

$$p_{AB}^1 < p_{AB}^* < p_{AB}^2 \quad (24)$$

$$p_{XX}^{13} < p_{XX}^{24} < p_{XX}^{ally} < p_{XX}^{non-ally} \quad (25)$$

From (23) and (24), carrier 1 (the alliance partner) charges lower domestic fares than does carrier 2, the nonallied airline. Moreover, carrier 1's domestic fares lie below those charged in the symmetric cases considered above, while carrier 2's lie above the symmetric fares. With suitable relabeling, (23) and (24) apply to the domestic fares charged by carriers 3 and 4. The major implication of the asymmetric solution, however, is that the alliance partners charge *a lower international fare* than the nonallied carriers. In contrast to the domestic fares, however, both international fares are cheaper than those charged in the symmetric cases.⁶

The impact of economies of traffic density helps to explain the pattern of fares in the asymmetric case. As before, internalization of the subfare externality tends to reduce the alliance's international fare, pushing it below that charged by the nonallied carriers. The resulting gain in alliance traffic then leads, via economies of density, to lower marginal costs, which puts further downward pressure on the alliance fare. Conversely, the traffic loss for nonallied carriers raises marginal cost, leading to upward pressure on the nonallied fare. These marginal-cost changes in turn affect domestic pricing, allowing the alliance partners to cut fares while forcing nonallied carriers to raise their domestic fares. The resulting gains and losses in domestic traffic further reinforce the marginal-cost advantage of the alliance partners. Confirming the role of economies of density, it may be shown the nonallied/allied fare difference is an increasing function of θ in each market. Thus, stronger economies of density lead to a greater fare advantage for the alliance partners.

3. The Effect of Alliances on Interline Fares

3.1. Overview

The first hypothesis we test, which follows directly from the theoretical results above, is that alliance partners charge lower interline fares than nonallied carriers. This hypothesis is

tested by estimating the following reduced-form regression:

$$\text{fare} = F(\text{distance; demand; competition; airline-specific effects; alliance dummy}). \quad (26)$$

Under the maintained hypothesis, the alliance dummy variable has a negative coefficient.

To understand this regression, an explanation of the nature of the data is helpful. The data are drawn from the U.S. Department of Transportation's Origin and Destination Survey, Databank 1A (DB1A), for the third quarter of 1997. The DB1A is based on a 10 percent quarterly sample of all airline tickets where one or more route segments is flown on an American carrier. Each record in the DB1A shows the itinerary (route and carriers used), the fare, the distance travelled, and number of passengers observed on the given itinerary at the given fare. The data may contain multiple records for a given itinerary, each showing a different observed fare. These records can be collapsed into a single observation by computing a passenger-weighted average of the multiple fares, which is then assigned to the given itinerary.

Each carrier/route itinerary then generates a single record, and these records constitute the observations for the fare regression in (26). The observation indicates the fare, distance of the trip, and the carriers used, as required in (26). Variables measuring the level of demand and the extent of competition in the itinerary's city-pair market are computed, as discussed below. Then, supplementary information, also discussed below, is used to construct an alliance dummy variable, which assumes the value one if the carriers on the itinerary are allied. While the interline fare depends on the identities of the individual carriers used for trip, this variable captures an additional impact, analogous to an interaction effect, which emerges if the carriers are allied. Assessing the direction of this impact is the goal of the empirical work.

It should be noted that the theoretical model implies that the interline fare should be affected by traffic densities on the various route segments used for the trip, with higher densities leading to a lower fare. Although Brueckner and Spiller (1994) included densities in their empirical model of domestic fare determination, this approach is unworkable in an international context. The reason is that traffic-density data for foreign carriers are unavailable.

3.2. Construction of data sets

The raw data set contains 2.7 million observations, and we restrict it in various ways to make the empirical work tractable. First, we eliminate purely domestic U.S. traffic, as well as itineraries with surface segments. Second, we eliminate itineraries with fares below \$100, which may represent trips purchased with frequent flier miles or made by airline employees at significantly reduced fares. Third, the itinerary must represent round-trip travel, with the same starting and ending airports, and it can have only one trip break. Finally, itineraries involving first class travel are removed. The data set resulting from these conditions is denoted the “restricted” data set.⁷

In order to study the effects of alliances on interline fares, further limitations are imposed on the restricted data set. First, because two-segment itineraries, which are nonstop in both directions, do not involve standard interline trips, such itineraries are eliminated.⁸ In addition, to reduce the complexity of the trips considered, itineraries using more than 2 carriers are dropped. For the same reason, highly unusual and circuitous trips are eliminated by restricting attention to itineraries with 3 or fewer segments in each direction (for a maximum of 6 segments over the whole trip). 82 percent of the records in the restricted data set met this condition. The final step is to compute passenger-weighted fares for each itinerary, as discussed above. The data set resulting from these additional restrictions is denoted the “interline” data set, and it contains 46,620 observations, representing 16,765 distinct city-pair markets.

Because the DB1A is restricted to tickets with at least one segment flown on an American carrier, the interline data set provides an incomplete picture of foreign carrier service. Foreign carriers offering online service in a market, or foreign carrier pairs offering interline service, are unobserved. This can understate the extent of competition in a market, whose measurement is discussed in detail below.

To see whether this problem affects our results, we explore the effect of an additional restriction on the interline data set. This restriction relies on the following observation: any itinerary whose origin or destination is a U.S. airport that has no foreign carrier enplanements must contain at least one route segment flown on a U.S. carrier, making it observable in the data. Endpoints that fit this criterion are simply U.S. behind-the-gateway cities. Thus, if we

focus on a data set restricted to itineraries which have a U.S. behind-the-gateway city as origin or destination, the problem of unobserved foreign competition is mitigated. The “behind-the-gateway” (or BTG) data set is constructed using information from the Official Airline Guide (OAG), which allows identification of U.S. airports without foreign carrier service.⁹ The BTG data set contains 11,694 observations, representing 6,917 city-pair markets.

We also construct two additional data sets pertaining only to U.S. passengers. The first is the “U.S.-origins” data set, which includes itineraries from the interline data set that originate at U.S. airports. This data set contains 26,039 observations. Since this data set also suffers from potential mismeasurement of competition, we also construct the “U.S.-origins/BTG” data set, which further restricts the sample to itineraries with origins at U.S. behind-the-gateway cities. The results from both the data sets complement those from the interline and BTG data sets by showing the effect of alliances on U.S. passengers, including those originating at smaller cities.

3.3. Constructing the alliance dummy and demand variables

To determine which carriers operate alliances, we use the list of current codesharing agreements from the U.S. Department of Transportation, U.S. Air Carrier Licensing Division. This list distinguishes alliances that are still in operation from those that have become dormant. We observe 50 alliance pairs in the restricted data set, as shown in Table 1, although four alliances dominate. Northwest-KLM carried 23.0 percent of the sampled alliance passengers, American-Canadian Airlines 14.9 percent, United-Lufthansa 12.5 percent and United-Air Canada 9.4 percent.

It should be noted that even when two carriers have a codesharing agreement, some of their interline itineraries may not involve codeshare flights. However, since codeshare and noncodeshare flights are not separately identified in the DB1A, it is impossible to measure the separate impact of codesharing on fares.¹⁰ All that can be measured is the overall impact of the alliance, as reflected in the existence of a codesharing agreement.

The first variable used to capture demand effects in a city-pair market is city population. U.S. city populations come from Census data, while foreign city populations are drawn from several sources.¹¹ Whenever possible, city populations represent the entire metropolitan area.

This information yields populations for the itinerary’s origin and destination cities, which are denoted POP_ORI and POP_DES, respectively. Following Brueckner and Spiller (1994), we assume the demand for air travel in a city-pair market depends on the market’s “population potential,” as measured by the geometric mean (in millions) of POP_ORI and POP_DES. This variable is denoted POP_POT.

Demand for travel in a city-pair market also depends on passenger incomes. Since the observed itineraries indicate the direction of travel in each market, the appropriate demand variable is a measure of income at the origin city. Thus, 1996 per capita for the origin city, denote INC, is included as a demand variable for observations in both data sets based on U.S. origins.¹² Since foreign city incomes are not readily available, this income variable does not appear in the regressions based on the interline and BTG data sets, which include foreign origins.

3.4. Measuring Competition

The most significant source of competition in a city-pair market is online service, which occurs when an itinerary is flown on a single carrier. In addition, competition can also come from interline flights, operated either by alliances or nonallied carriers. For reasons that are explained below, we measure nonallied interline competition separately from online and alliance competition. Online and alliance competition are measured in one variable, but we construct that variable in several different ways to account for various potential problems.

To construct competition measures, we must first determine which airlines offer service in a market, either online or interline. One possibility would be to consult the schedules provided by the Official Airline Guide. However, because the OAG’s coverage of connecting flights is incomplete, and because the number of city-pairs in the data is huge (over 23,000), use of this source is ruled out. Instead, our approach is to measure competition using the itineraries observed in the restricted data set.

By decomposing the information in each observed itinerary, we can generate the city-pair combinations served by each airline. For example, suppose the outbound portion of an itinerary shows travel from Indianapolis (IND) to Detroit (DTW) and then to Amsterdam (AMS), all on Northwest, with a final segment to Paris (CDG) flown on KLM. We could then infer that

Northwest provides online service in the IND-AMS and DTW-AMS city-pair markets, while KLM provides online service in the AMS-CDG market. The reasoning is that any subset of an itinerary flown on a single carrier must also represent a legitimate trip on that carrier.

The same methodology is applied to capture interline service (both alliance and non-alliance). In this case, subset itineraries flown on two carriers are counted as interline service. In the above example, Northwest and KLM are counted as providing interline service in the IND-CDG and DTW-CDG markets. By repeating this exercise for each itinerary in the restricted data set, we build a picture of available online and interline service in all city-pair markets.

To measure the extent of competition from this information, we cannot simply count the different types of online and interline service (both alliance and nonalliance) available in a market. This will lead to various kinds of double counting. To see the problem, consider the Houston-Toronto market, and suppose that two carriers, United and Air Canada, offer online service. One might be tempted to call this market a duopoly, but because United and Air Canada are alliance partners, there is no reason to believe that they compete in a market where they both offer service. As a result, this market would be better described as a monopoly. The example becomes more complex if we suppose Continental also offers online service. Continental and Air Canada are also alliance partners but United and Continental are not. Is it more accurate to describe this market as a monopoly or a duopoly?

To handle this problem, we construct three measures of online and alliance competition for a city-pair market. The first, which ignores the problem above, is simply a count of the number of airlines providing online service plus the number of alliances offering interline service. An alliance is not counted if one of its partners also offers online service in the market. This avoids the double counting that arises when an airline offers online service in a market but also offers interline service with an alliance partner. We call the resulting variable TOTFIRM.

The other two methods are designed to address the problem illustrated above and represent lower and upper bounds on the number of online and alliance competitors. The lower bound assumes that all markets like the one above behave as monopolies. Generally, this measure counts the number of airlines offering online service in the market but assumes that when two nonallied airlines in a market are both allied with a third carrier also in the market, then all

three carriers behave as one competitor. From the example above, United and Continental are not allied but are connected through their alliances with Air Canada. Thus, all three carriers are assumed to behave as one firm. This method generates a “lower bound” because it measures the fewest number of competitors one could reasonably believe operates in the market. The resulting competition measure is denoted COM_LOW.

The upper bound on competition also counts the number of airlines offering online service but assumes that each alliance pair (again where both airlines offer online service) behaves as one firm in the market.¹³ In the example above, this method considers the market a duopoly, served by the United-Air Canada alliance and the Continental-Air Canada alliance. This measure is called the “upper bound” because it assumes full competition between carriers that are not allied themselves. We call this variable COM_HIGH.

As with the TOTFIRM variable, we add interline alliance competitors, where neither airline offers online service, to both COM_LOW and COM_HIGH. The following relationship between the three measures always holds: $TOTFIRM \geq HIGH_COM \geq LOW_COM$. The true amount of competition in each market is likely to lie within these bounds.

As mentioned earlier, we create a separate variable to measure the effect of nonallied interline competition. We expect additional nonallied competitors to reduce fares, but excessive interlining may occur in markets with more restrictive bilateral agreements and thus, more monopolistic behind-the-gateway routes and higher fares. To compute this variable, we begin by eliminating all interline itineraries in a city-pair market where a component carrier operates either online or alliance service. Then we count each remaining carrier as $\frac{1}{2}$ of a competitor. The reason for constructing the variable in this fashion rather than simply counting nonallied itineraries is that it is not uncommon to see an airline fly one portion of a nonallied trip, while several different carriers fly the other portion. We call the resulting variable COM_INT.

3.5. Summary statistics

Table 2 provides a list of the variables used in the analysis, and Tables 3–5 present the summary statistics for three data sets: interline, BTG, and U.S.-origins (U.S.-origins/BTG statistics are omitted). For each data set, the figures are broken into alliance and non-alliance categories.

AVGFARE is the average passenger-weighted fare for the itinerary. In the regressions, we use its log transformation, LNFARE. The dummy variable ALLY indicates whether the carriers providing the itinerary's service are alliance partners. Tables 3–5 show that between 39 and 43 percent of the itineraries are flown on allied carriers across the three data sets. In addition, the Table reveals that alliance fares are lower on average than the fares charged by nonallied carriers, as predicted by our hypothesis. Alliance fares are 19 percent lower on average in the interline data set, although the average is only 15 percent lower in the BTG data set.

TOT_DIST is the total distance flown on the itinerary. Two itineraries representing the same city-pair market will have different values for this variable if one involves a more circuitous route. In the regressions, distance is log transformed, and the resulting variable is called LNDIST. Because this variable captures the higher operating cost for longer flights, we expect it to have a positive coefficient in the regressions. In the summary statistics, it is worth noting that the average distance traveled on alliance itineraries, which have lower average fares, is actually longer than on non-alliance itineraries.

A related variable, COUP, equals the number of ticket coupons and thus the number of flight segments in the itinerary.¹⁴ From the cost side, more coupon segments should result in higher costs and fares, but from the demand side, consumers should generally need compensation for flights with more coupon segments. In reduced form, the effect of this variable is uncertain.

As discussed above, the demand variables POP_POT and INC are expected to have positive coefficients in the regressions. Tables 3–5 show that mean values for POP_POT is smallest in the BTG data set, indicating that cities in behind-the-gateway markets are smaller on average than those in the larger data sets. In addition, U.S.-origins/BTG data has a smaller mean value of INC (not shown) than the U.S.-origins data set, indicating that U.S. behind-the-gateway origins are less affluent than the average U.S. origin city.

Competition in the market is measured by the variables whose creation was discussed above, and their coefficients are expected to be negative. TOTFIRM is the number of carriers operating in the market, and COM_INT is the number of nonallied interlining partners. COM_LOW and COM_HIGH are the lower and upper bounds on the number of online and alliance competi-

tors in the market. In the interline and U.S.-origins data set, the average itinerary's city-pair market has 3–3.5 online/alliance competitors and 1.2–1.3 nonallied interline competitors. In the BTG data set, the number of alliance/online competitors is substantially lower on average, at 1.8, while nonallied interline competition is only slightly lower, with the average number of competitors equal to 1.0. Thus, smaller markets enjoy less competition than larger markets.

In the BTG data set, TOTFIRM, COM_LOW and COM_HIGH are equal for 92 percent of the observations. In an additional 7 percent of the data, COM_LOW and COM_HIGH are equal but different from TOTFIRM. Thus, there is very little difference between our competition measures in the behind-the-gateway markets. By contrast, in the full data set, all three measures are equal for only 70 percent of the observations, with the bounds equal but different from TOTFIRM in an additional 18 percent. This difference arises because the BTG data set excludes the larger markets, where there is greater chance of alliance partners operating together.

The remaining variables are region dummies indicating the region of the destination airport. Because non-U.S. origins are included in the interline and BTG data sets, a U.S. destination dummy is included. However, since the remaining data sets are restricted to U.S. origins, the U.S. destination dummy is omitted. In all regressions, Europe is the omitted region.

Although omitted from the summary statistics, the regressions also include airline-specific effects. Including such variables allows us to separate the alliance effect from the effect of carrier characteristics such as labor cost, which influence fares. We could simply include a dummy variable for each of the 227 airlines in the data, but instead we use a measure more closely related to variable costs. The measure equals the interaction between the carrier dummies and distance flown. For example, suppose the itinerary involves 3000 miles of travel on United and 500 miles of travel on Lufthansa. Our method sets United's variable equal to $\ln(3000)$ and Lufthansa's to $\ln(500)$, with all other airline variables set equal to zero (the omitted carrier is American). Because United and Lufthansa are allied, ALLY would equal one for this observation. It should be noted that under this specification, the alliance effect is assumed to be independent of the identities of the allied carriers.

3.6. Regression Results

The model we estimate is of the following form:

$$\begin{aligned} \text{LNFARE}_{ij} = & \beta_0 + \beta_1 \text{ALLY}_{ij} + \beta_2 \text{LNDIST}_{ij} + \beta_3 \text{COUP}_{ij} \\ & + \beta_4 \text{Online/Alliance Competition}_j + \beta_5 \text{COM_INT}_j \\ & + \beta_6 \text{POP_POT}_j + \beta_7 \text{Regions}_j + \beta_8 \text{Airline Effects}_{ij} + (v_j + \epsilon_{ij}). \end{aligned} \quad (27)$$

“Online/Alliance Competition” denotes one of the variables TOTFIRM, COM_HIGH or COM_LOW, “Regions” refers to the destination dummies, and “Airline effects” denotes the airline variables. In (27), j refers to the city-pair market, i refers to the itinerary within that market, and ϵ_{ij} is the error term. Note that while the competition, population potential and region dummy variables are constant across the itineraries within a given market, the remaining variables depend on the itinerary. Note also that this specification allows for a fixed effect, v_j , in each city-pair market, which captures unobservable attributes of the market.

We begin with simple OLS regressions on the interline data set, ignoring the possibility of market fixed effects. The results appear in Table 6. With the alliance dummy as the only regressor, the estimates show that alliance partners charge fares 19.5 percent lower than those charged by nonallied carriers. In column (2), we add distance and coupons to the regression, and the alliance effect increases to 24.7 percent. This confirms what we know from the summary statistics, namely that alliance itineraries are longer on average. As we add the demand, competition (COM_HIGH and COM_INT), and region dummy variables in columns (3)-(5), the conclusion changes very little: the alliance dummy coefficient is negative and highly significant, indicating that alliance fares are lower by 24-26 percent. Finally, in column (6), we introduce carrier-specific effects, and the coefficient increases to 26.4 percent, suggesting that higher-cost carriers are more likely to form alliances.

In the previous regressions, we arbitrarily chose COM_HIGH to measure the effect of online and alliance competition. In columns (6)–(8), we consider the effects of using the other measures of competition, TOTFIRM and COM_LOW. All three measures have roughly the same coefficient, and the choice among them has no appreciable effect on the alliance

dummy. Each additional online/alliance competitor lowers the interline fare by approximately 1.5 percent, and additional non-alliance interline competitors generate about half the impact, at 0.8 percent.

The coefficients on the other regressors are as expected. Greater distance increases the fare, but the consumer must be compensated for additional coupon segments with a lower fare. Greater market size, as measured by the population potential variable, leads to a higher fare by creating stronger demand.¹⁵

To address the problem of potential mismeasurement of competition in the interline data set, we run the same regressions on the BTG data set, where all competition is potentially observable. The results are shown in Table 7. In column (1), where the alliance dummy is the only regressor, the effect of alliances is smaller than in the previous results, with alliance fares only 13 percent lower than those charged by nonallied carriers. However, addition of distance and coupon segments to the regression has a large impact on the alliance effect, raising it to 21 percent. In column (6), where all the regressors are used (including airline-specific effects), the effect jumps to 26 percent. In all specifications, the alliance coefficient is highly significant. Since the coefficients in column 6 of both Tables are virtually the same, the estimated fare effect of alliances is uninfluenced by any mismeasurement of competition.

As before, columns (6)–(8) of Table 7 show that the method used for counting competitors has little effect on the BTG results. Additional online/alliance competitors still significantly decrease fares, as do non-alliance interline competitors.

The results for the U.S.-origins data set appear in Table 8, and those for the U.S.-origins/BTG data set are shown in Table 9. Although the results are qualitatively similar to those in Tables 6 and 7, some coefficient magnitudes differ. At 23 percent, the alliance coefficient from column (6) is slightly smaller in the U.S.-origins data set than in the earlier Tables. The coefficient is even smaller, at 19 percent, in the U.S.-origins/BTG results. These results suggest that the benefits from alliances are somewhat lower for U.S. than for foreign passengers, with residents of smaller U.S. cities receiving the lowest benefits (these are nevertheless still appreciable). The same conclusion emerges from a different specification using the larger data sets, where the alliance coefficient is interacted with a U.S.-origin dummy variable.

Although these results are not reported, the interaction coefficient is positive and significant in both the interline and BTG data sets, indicating a weaker alliance effect for U.S. passengers.

Being restricted to U.S. origins, the regressions in Tables 8 and 9 also include the origin city's per capita income, INC, which has a positive and significant coefficient, as expected. The remaining coefficients are also as expected, although the coefficient of COUP is positive and significant in Table 9. This coefficient captures a mixture of supply and demand responses, and the estimate suggests that in the U.S.-origins/BTG data set, the additional costs of extra stops are outweighing consumers' preferences for fewer stops. Note also that in the U.S.-origins results in Table 8, the effects of online/alliance competition are weaker, and the effects of interline competition stronger, than in Table 6's interline results. A similar comparison applies to the BTG estimates in Tables 7 and 9.

To this point, we have ignored market-specific effects. If they are significant, then the OLS coefficients will be biased. The traditional solution to this problem is fixed-effects estimation, but such estimation requires at least two observations for each city-pair market. In the interline data set, approximately 71 percent of the data fit this criterion, but only 46 percent of the BTG observations do. Estimates for the fixed-effect specification for the different data sets are found in column (9) of their respective Tables. The alliance dummy coefficient in these markets is still negative, significant, and of roughly the same magnitude as before, ranging between 21 and 28 percent in the four data sets. Although the joint tests for market-specific effects are significant at the 99 percent level, controlling for these effects appears to have little impact on the alliance coefficient.¹⁶

Since the fixed-effects estimates ignore the smaller, more monopolistic markets, the coefficients may be misleading. To retain these small markets, the OLS regressions are repeated on a modified data set where we randomly select one observation from each city-pair market that has multiple observations. This data set eliminates the market-specific effect by construction, but underweights the denser markets. The results are found in column (10) of each Table. The alliance dummy coefficient lies between 18 and 23 percent and is again highly significant. These effects are in line with the OLS and fixed-effect estimates, though somewhat smaller in magnitude. The remaining coefficients are also similar to those in previous regressions.

The consistency of our results across specifications and estimation techniques is remarkable. The results suggest that for interline itineraries, the fares charged by alliance partners are 18 to 28 percent lower than those of nonallied carriers.

4. The Effect of Alliances on Gateway-to-Gateway Fares

4.1. Overview

Although the gateway-to-gateway market was suppressed in the model of section 2, the impact of alliances in such markets is an important concern of government regulators, as discussed in the introduction. The concern is that when two carriers both provide service in a given gateway-to-gateway market, an alliance between them will be anticompetitive, leading to higher fares in the market. This prediction is confirmed by the analysis of Brueckner (1997), as explained above.

Given the importance of gateway-to-gateway markets, this section provides empirical evidence on the effect of alliances in such markets. The empirical strategy is to focus on nonstop service between gateway cities, running a regression that relates fares to distance, demand, airline specific effects, and competition variables. The first competition variable is the previous TOTFIRM measure, which counts the total number of carriers in the market. The second variable measures the loss of competition relative to this total due to the presence of alliances among carriers in the market. The hypothesis is that the reduction in competition associated with alliances leads to higher fares, implying a positive coefficient for the loss measure.

4.2. Gateway-to-gateway data sets

To conduct this analysis, we construct a data set of gateway-to-gateway itineraries from our original restricted data set, keeping only round-trip, two segment itineraries that are operated by one carrier. This “GTG” data set contains 1,313 observations, representing 663 city-pair markets. Like the interline data sets, the GTG data set also suffers from unobserved city-pair effects. However, because the parameters of interest in this analysis are constant within a city-pair market, fixed-effects estimation is not an option. Instead, we create two additional data sets that are distinguished by their treatment of market-specific effects. For the first, we take the same approach as in the interline analysis and select one observation at random

from city-pair markets with multiple observations. The other data set solves the problem by aggregating the data to the city-pair level instead of preserving individual itineraries. A passenger-weighted average fare is computed, but in this case, we average across the itineraries of all the carriers serving the market. We refer to these data as the city-pair data set. Since only 25 percent of the observations occur in markets with multiple observations, the reduced data sets are not much smaller than the GTG data set, containing 1,137 observations.

4.3. Variables and summary statistics

Because we restrict our attention to nonstop service, we use the OAG to determine which carriers operate in each city-pair market.¹⁷ Thus, the mismeasurement issue that arose in the interline data is not present. As noted above, the variable measuring the total number of airlines operating in a market is TOTFIRM. To measure lost competition from alliances, we define two new variables, LOSS_HI and LOSS_LO, which are simply the difference between TOTFIRM and COM_HIGH and COM_LOW, respectively. The loss measure gives the number of airlines in the market that do not provide competition because of the presence of alliances. For example, suppose a market is served by United, American and Air Canada. TOTFIRM in this market equals three. However, because United and Air Canada are allied, both COM_HIGH and COM_LOW equal two. Thus, the loss measures for this market are both equal to one, meaning one of the airlines providing service in the market is not considered a competitor because of alliances.

The loss measures differ from zero for only 10 percent of the observations, implying that alliance partners rarely both operate flights in the same gateway-to-gateway market. This number could be small for two reasons. On the one hand, airlines may choose partners where there is little network overlap, in which case anticompetitive effects seldom arise in gateway-to-gateway markets. Alternatively, after alliances are formed, one airline may stop operating flights in the market. In this case, the anticompetitive effects may be large but missed by our analysis.

Since itineraries in the gateway data sets are flown on the same route by a single airline, we measure airline specific effects with simple airline dummies (the distance interaction is not needed). We only observe 22 airlines in the gateway-to-gateway markets, and as before,

American Airlines is omitted from the regressions. In the city-pair data set, where itineraries are aggregated across airlines, the dummy is split between the operating carriers on a passenger-weighted basis. For example, suppose United and Continental are the only airlines operating in a particular gateway-to-gateway market, with 60 passengers observed using United and 40 using Continental. United’s dummy for this city pair market would then be set equal to 0.6 and Continental’s to 0.4.

The summary statistics for the gateway data sets are presented in Table 10. As can be seen, the gateway-to-gateway itineraries have a different character than those in the interline data set. The average fare is around \$826, and the average distance is only 5,700 miles. The average number of carriers operating nonstop flights in the market is 2.14 but drops to 1.90 when city-pairs are counted only once (in both the random-observation and city-pair data sets). This indicates, not surprisingly, that we observe multiple itineraries in markets with more competition. The mean values of the loss measures, LOSS_LO and LOSS_HI, are 0.21 and 0.16 respectively, with the means slightly smaller in the city-pair level data.

4.4. Regression Results

The model we estimate for the gateway-to-gateway markets is

$$\begin{aligned} \text{LNFARE}_{ij} = & \beta_0 + \beta_1 \text{TOTFIRM}_j + \beta_2 \text{Loss measure}_j + \beta_3 \text{LNDIST}_j \\ & + \beta_4 \text{POP_POT}_j + \beta_5 \text{Regions}_j + \beta_6 \text{Airline Dummies}_{ij} + (v_j + \epsilon_{ij}). \end{aligned} \quad (28)$$

“Loss measure” refers to either LOSS_HI or LOSS_LO. Observe that, with the exception of the airline dummies, all the variables in (28) are constant across different itineraries i within a given city-pair market j .

Notice that by using TOTFIRM along with the loss measure, we can distinguish the effect of adding a nonallied airline to the market from the effect of adding a carrier that is allied with another airline in the market. The effect of adding a nonallied airline is given by the coefficient of TOTFIRM, while the effect of adding an alliance partner is equal to the sum of the coefficients of TOTFIRM and the loss measure. We expect the coefficients on TOTFIRM and the loss measure to be negative and positive, respectively. In addition, the sum of the

coefficients should be close to zero, since adding an alliance partner of an existing carrier should have little effect on the fare in the market. We test this prediction using both LOSS_HI and LOSS_LO.

The gateway-to-gateway regression results are presented in Table 10, with two columns devoted to each of the data sets. In all specifications, we meet only two of our three expectations. In particular, the coefficient on TOTFIRM is negative and significant, with its magnitude indicating that an extra competitor reduces fares by about 5 percent. In addition, the sum of the coefficients on TOTFIRM and the loss measure is not significantly different from zero. However, the coefficient on the loss measure itself, while negative and of the proper 4-6 percent magnitude, is not significantly different from zero at the 5 percent confidence level. This is true regardless of which measure is used. The other coefficients have the expected signs and magnitudes.

One drawback to these regressions is that destinations for many itineraries are resort islands like Tahiti and the Bahamas, which are not gateways in the usual sense, lacking behind-the-gateway endpoints. Since such markets are not a particular concern of regulators, it is appropriate to delete them and evaluate the effect on the results. Dropping itineraries that start or finish in the Caribbean or on any south Pacific island reduces the data set by 23 percent. However, the resulting estimates show that this restriction of the sample has little effect on the magnitude and statistical significance of the loss coefficients.

Another possibility is that our model misspecifies the effect of competition by assuming that an additional competitor has the same fare impact regardless of the level of competition. However, if we allow the marginal impact of the second, third and additional competitors beyond three to differ, the performance of the loss measure is unaffected.

These additional results confirm the absence of a statistically-significant effect from lost competition due to alliances in the gateway-to-gateway markets. While such an effect may not be operative in actuality, it is more likely the effect is truly present but that the relatively small GTG sample sizes, along with the sparsity of observations where the loss measures are nonzero, rules out precise estimation of its magnitude.

4.5. Simulation analysis of the AA/BA alliance

The estimated effects of alliances on interline and gateway-to-gateway fares can be used to shed light on a current regulatory issue: the welfare impact of the proposed alliance between American Airlines and British Airways. Welfare gains for interline passengers are computed by assuming that fares fall by 26 percent in each interline market, as in column 6 of Table 6. Combining this figure with a value for the price elasticity of demand, we can compute the increase in consumer surplus for each AA/BA interline itinerary observed in the interline data set. Next, using the point estimate in Table 11 but ignoring its lack of statistical significance, we assume that fares will rise by 5 percent in each gateway-to-gateway market currently served by both AA and BA following an alliance. Then, using the same price elasticity value, we compute the loss in consumer surplus in these markets. The net change in surplus gives an estimate of the welfare impact of the alliance.

The interline data set contains 605 AA/BA itineraries, with a total of 809 sampled passengers. This corresponds to 8090 actual passengers, given the DB1A's 10 percent sampling rate. Assuming a price elasticity of -1 , the 26 percent fare decline raises traffic to 10,266 passengers and generates an aggregate consumer surplus gain of \$4,964,000. Using the larger elasticity estimate of -2.5 from Brueckner and Spiller (1994), traffic rises to 13,429 passengers, and the surplus gain is \$5,833,000. With the still-larger elasticity of -4 , traffic rises to 16,633 passengers, and the surplus gain is \$6,701,000. These results are summarized in Table 12.

AA and BA compete in six gateway-to-gateway markets, and AA carries 5,309 sampled passengers in these markets, corresponding to 53,090 actual passengers. We assume that BA traffic, which is not observed in the data, is identical to AA's, and that BA passengers experience the same fare effects. With an elasticity of -1 , the 5 percent fare increase caused by the alliance reduces AA's gateway-to-gateway traffic to 50,436 and generates a surplus loss of \$2,631,000, for a total loss across both carriers of \$5,263,000. When the elasticity equals -2.5 , AA's traffic falls to 46,454, and total surplus declines by \$5,060,000. With an elasticity of -4 , AA's traffic falls to 42,472, and total surplus declines by \$4,858,000.

These values, which are summarized in Table 12, generate net consumer surplus changes $-\$298,000$, $\$772,000$, and $\$1,843,000$ when the price elasticities are -1 , -2.5 , and -4 , respec-

tively. These numbers suggest that the welfare impact of the AA/BA alliance on passengers is, at worst, slightly negative, with benefits emerging if demand is relatively elastic.

Fare effects that are harder to quantify, however, will be felt by other passengers in the event of an AA/BA alliance, and these must be considered. The first affected group consists of AA passengers who currently make interline trips using carriers other than BA. In the event of an alliance, these passengers would presumably switch to BA for the non-AA portion of their trip, enjoying the fare reduction associated with alliance travel. Although, in principle, the resulting gain could be computed using our data, various complications would arise.¹⁸

Other passengers would experience indirect benefits through changes in the extent of competition following an AA/BA alliance. Since alliance competition is more effective than interline competition (recall that COM_INT has a smaller coefficient than the other competition variables), other carriers' interline fares will fall slightly in response to an AA/BA alliance. Conversely, the loss of competition will lead to higher fares for non-AA/BA passengers in the gateway-to-gateway markets. An obstacle to the computation of both effects is the absence of foreign-carrier traffic information. For example, since we do not observe Virgin Airways' U.S.-U.K. traffic, we cannot compute the surplus loss for its passengers resulting from lower gateway-to-gateway competition.

It is likely that if these additional effects could be reliably computed, their net impact on consumer surplus would be positive. First, since AA and BA control the majority of U.S.-U.K. traffic, the omitted gateway-to-gateway losses would be relatively small. In addition, the benefits for current AA/non-BA interline passengers would be substantial given the large volume of such traffic. Therefore, it is likely that the figures in Table 12 would be adjusted upward, leading to an unambiguous welfare gain from the AA/BA alliance. This conclusion would be reinforced if passenger benefits from the greater convenience of alliance travel, which have not been measured, could be added to the calculation.¹⁹

Despite the existence of this probable welfare gain, regulators are properly concerned about the anticompetitive effect of the proposed AA/BA alliance. Making approval of the alliance contingent on entry of new competitors in the gateway-to-gateway markets, as has been proposed, is sound regulatory policy.

5. Conclusion

This paper provides strong evidence on the effect of international airline alliances on fares. Our main finding is that alliance partners charge interline fares that are 18–28 percent below those charged by nonallied carriers. According to the theoretical model, this fare reduction arises from two sources. First, joint pricing internalizes the negative externalities from the uncoordinated choice of subfares, leading to lower fares. This fare reduction then stimulates traffic, which in turn lowers marginal costs via economies of traffic density, leading to further downward pressure on fares. Our results, which confirm this prediction, are extremely robust to changes in the measurement of competition and the handling of market-specific effects.

The evidence of an anticompetitive alliance effect in the gateway-to-gateway markets is not as clear. While our point estimates show that an alliance between two previously-competitive carriers would raise fares by 4–6 percent, this effect is not statistically significant. This finding, however, may be due to data limitations, which prevent precise measurement of the alliance effect. In addition, if alliances lead to changes in the route structures of the partners, anticompetitive effects may exist that we simply cannot measure in cross-section data. For example, if one partner exits the gateway-to-gateway market after an alliance is formed, fares in the market will rise by approximately 5 percent, and this will go unnoticed in our analysis. To address this possibility, future research could analyze the changes in network structure that occur when an alliance is formed.

Further work could also focus on estimating a structural model to disentangle the supply and demand effects of alliances. With closer gates, better-timed connections and merged frequent flier programs, alliances are meant to provide the consumer with a better product than the traditional interline trip. If this is so, the lower fares we observe are net of any increases due to the effect of stronger demand. In principle, this demand effect could be isolated by a structural model.

The newly-proposed U.S. domestic alliances have gained a great deal of press attention lately, and our research may also shed light on their competitive effects.²⁰ In contrast to international alliance partners, whose networks overlap only on gateway-to-gateway routes, the competition between the proposed domestic alliance partners is more widespread. While

they compete directly only on a small number of nonstop routes, competition is extensive in markets involving small and medium-size cities, which are served through different hub airports. Given this greater extent of competition, the anticompetitive impact of domestic alliances is likely to be more significant than in the international case.

References

- BRUECKNER, J.K. AND P.T. SPILLER, 1991, Competition and mergers in airline networks, *International Journal of Industrial Organization* 9, 323-342.
- BRUECKNER, J.K. AND P.T. SPILLER, 1994, Economies of traffic density in the deregulated airline industry, *Journal of Law and Economics* 37, 379-415.
- BRUECKNER, J.K., 1997, The economics of international codesharing: An analysis of airline alliances, Unpublished paper, University of Illinois at Urbana-Champaign.
- CAVES, D.W., L.R. CHRISTENSEN AND M.W. TRETHERWAY, 1984, Economies of density versus economies of scale: Why trunk and local service costs differ, *Rand Journal of Economics* 15, 471-489.
- DOGANIS, R., 1985, Flying off course: The economics of international airlines (George Allen & Unwin, London).
- GELLMAN RESEARCH ASSOCIATES, 1994, A study of international codesharing (Office of International Aviation, U.S. Department of Transportation, Washington, D.C.).
- MORRISON, S.A. AND C. WINSTON, 1995, The evolution of the airline industry (Brookings Institution, Washington, D.C.).
- OUM, T.H., J.-H. PARK AND A. ZHANG, 1996, The effects of airline codesharing agreements on firm conduct and international air fares, *Journal of Transport Economics and Policy* 30, 187-202.
- PARK, J.-H., 1997, The effect of airline alliances on markets and economic welfare, *Transportation Research-E* 33, 181-195.
- PARK, J.-H. AND A. ZHANG, 1998, An empirical analysis of global airline alliances: Cases in the north Atlantic markets, Unpublished paper, City University of Hong Kong.
- U.S. GENERAL ACCOUNTING OFFICE, 1995, Airline alliances produce benefits, but effect on competition is uncertain (General Accounting Office, Washington, D.C.).
- YOUSSEF, W. AND M. HANSEN, 1994, Consequences of strategic alliances between international airlines: The case of Swissair and SAS, *Transportation Research-A* 28, 415-431.
- _____ 1993, *Cities of the World* (Gale Research Inc, Chicago).

_____ 1997, *Official Airline Guide: Worldwide Addition* (Reed Travel Group, Oakbrook, IL).

_____ 1993, *Rand McNally World Atlas* (Rand McNally and Company, Chicago).

Footnotes

*The authors are respectively Professor of Economics and Ph.D. candidate in economics. We thank Dan Bernhardt for helpful discussions.

¹Although interline service is feasible in market AB, we assume that consumers use a single carrier, making an online trip.

²With constant marginal costs, pricing in the various markets is independent, with fares determined by each individual first-order condition. Since (4) and (5) do not involve the parameter ϕ , it follows that the fares in markets XH and AB are independent of ϕ . The first-order condition in the XX markets can be written

$$\phi D(p_{XX}, p_{XX})/D_1(p_{XX}, p_{XX}) + p_{XX} - \delta = 0,$$

where the carrier superscripts are removed given symmetry, and where the D subscript denotes the partial derivative with respect to the first argument. Differentiating this equation yields

$$\frac{\partial p_{XX}}{\partial \phi} = -D[(p_{XX}D_{11} + (1 + \phi)D_1 - 2\delta D_{11}) + (p_{XX}D_{12} + \phi D_2 - 2\delta D_{12})]^{-1}.$$

The expression inside the first set of parentheses must be negative by the second-order condition, while it can be shown that the expression inside the second set of parentheses must be negative for the Cournot equilibrium to be stable. Therefore, the derivative is positive, as claimed above.

³Note that the AB fare is less than twice the XH fare. This rules out fare arbitrage, where an AB passenger could reduce his outlay by purchasing separate AH and BH tickets.

⁴Observe that the international fare in the alliance case is exactly twice the XH domestic fare. This means that an international passenger is indifferent between purchasing an international ticket and two domestic tickets to make his trip.

⁵The second-order conditions for the allied and nonallied carriers are $\alpha - (7/2)\theta > 0$ and $\alpha - (5/2)\theta > 0$, with the first being more stringent (these are derived from the relevant Hessian matrices). If $\alpha - 4\theta < 0$ holds, then the second-order condition implies $(7/2)\theta < \alpha < 4\theta$. But noting that the q_{AB}^1 solution is given by $(3\alpha - 10\theta)/6(\alpha - 4\theta)$, it follows that $3\alpha - 10\theta < 0$ must then hold for q_{AB}^1 to be positive. However, this implies $\alpha < (10/3)\theta$, which contradicts the previous inequality given $10/3 < 7/2$. With $\alpha - 4\theta$ thus positive, inspection of the other traffic solutions generates a necessary condition ensuring that each

is positive. This condition is the stronger inequality $\alpha > (14/3)\theta$. Marginal costs should also be positive, indicating that spoke traffic densities lie on the upward-sloping part of the quadratic cost function. The condition for positive marginal costs is $\alpha > 16\theta(1-2\theta)/(4-9\theta)$, which also ensures that all fares are positive. It can be shown that this constraint is more restrictive than the previous inequality for large θ 's, and less restrictive for small θ 's.

⁶It can be shown that, as discussed in previous footnotes, there are no possibilities for fare arbitrage in the asymmetric case.

⁷In most cases, small domestic feeder airlines share the code of the major carrier with whom they operate (for example, American Eagle is listed as American Airlines). Occasionally, however, the feeder carriers are listed under their own airline code. In such cases, the feeder code is changed to the parent carrier's code, although leaving the original code has no significant effect on the results.

⁸Use of different carriers on the inbound and outbound segments would represent interline travel, but this type of itinerary is not our focus.

⁹An alternative approach is to use a slightly-outdated Federal Aviation Administration (FAA) list of airports that had foreign carrier enplanements (the year is 1996). The regression results under the two approaches are very similar.

¹⁰From 1998 onward, codeshare flights are identified in the DB1A. This change, however, occurred too late to be of use in the present study.

¹¹U.S. city populations are estimates as of July 1997. Since these estimates are not available for towns with populations under 10,000, 1990 census data are used instead and adjusted for the state's population growth from 1990-1996. The sources for foreign city populations are the *Rand McNally World Atlas*, *Cities of the World* encyclopedias, and the U.S. State Department Web Site.

¹²U.S. city incomes are drawn from 1996 data provided by the Bureau of Economic Analysis, Department of Commerce. For towns too small to be covered by this data source, county data are substituted.

¹³This calculation is complicated by the presence of "group alliances," which occur when airlines in a group are allied with every other member of the group. For example, Delta, Sabena and Swissair form a group alliance because each is allied with the other two. When all the airlines in a group alliance offer online service in the same market, they are counted as one firm.

- ¹⁴When a passenger continues on the same flight after an intermediate stop, the data do not show a segment break. In these cases, which are relatively rare, COUP understates the number of flight segments.
- ¹⁵The one anomalous result occurs in column (3), before we introduce population controls, where the coefficient on the COM_HIGH is positive. Since population and competitors are positively correlated but have different effects on fares, the competition measure is likely biased by the omission of population.
- ¹⁶The only unusual result occurs in the BTG data set, where the coefficient on LNDIST is not statistically different from zero. This outcome is natural, however, given that the estimate is based on variation in this variable across itineraries within a given city-pair market, which is slight.
- ¹⁷We count only those carriers listed in the OAG as providing direct service in the market. In some cases, this includes one-stop service where the continuing flight retains the same flight number. Carriers providing connecting one-stop service are not counted, although including such competitors has little effect on the results.
- ¹⁸While each AA/non-BA interline itinerary that has a European origin or destination could be converted to an AA/BA itinerary, a precise calculation of the fare impact would require a change in the routing so that the trip passes through BA's London-Heathrow hub, with consequent changes in the distance measure and in the distance-based airline specific effects. Given the need for such speculative assumptions, we chose not to carry out these calculations.
- ¹⁹A welfare gain from the alliance is consistent with the theoretical results of Brueckner (1997), who shows that an alliance leads to higher consumer surplus over most of the feasible parameter space in his model.
- ²⁰The alliances are between Northwest/Continental, United/Delta, and American/US Airways. The U.S. Department of Justice has recently challenged the Northwest/Continental alliance.