International Journal of Industrial Organization 30 (2012) 735–747

Contents lists available at SciVerse ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Strategic alliance as a response to the threat of entry: Evidence from airline codesharing☆

Christopher F. Goetz [a], Adam Hale Shapiro [b,*]

[a] University of Maryland, U.S. Census Bureau, United States
[b] Federal Reserve Bank of San Francisco, United States

ARTICLE INFO

Article history:
Received 9 September 2011
Received in revised form 25 June 2012
Accepted 13 August 2012
Available online 25 August 2012

Keywords:
Strategic alliances
Entry deterrence
Airlines
Codesharing

ABSTRACT

Strategic alliances are arrangements in which firms combine efforts and resources to jointly pursue a business objective while remaining separate entities. An example of such a practice is airline codesharing, in which allied carriers engage in the cooperative marketing of certain flights. We empirically test for the presence of competitive motives behind such alliances by identifying an incumbent airline's use of codesharing in response to the threat of future entry by a competitor. Using within-flight segment, fixed-effects regressions on panel data from 1998 to 2010, we estimate the impact of exogenous threats of entry on an airline's decision whether to codeshare with a partner on a specific segment. Estimates show that when an incumbent carrier's segment is threatened by a low-cost competitor it is approximately 25% more likely than average to be codeshared with its partner. Further tests show that this effect depends strongly upon the level of market share that the airline has on the segment in question. We interpret this as evidence of a strategic alliance being used to preemptively act in anticipation of future competition.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The formation of strategic alliances between firms has become widespread throughout many industries. A strategic alliance is defined as any voluntary partnership that represents neither a simple transactional relationship, nor a significant structural merging of the entities. The stated goals of such arrangements are typically to reduce transaction costs, share risk, integrate networks, or otherwise mutually add value or reduce cost, all of which are features that have been covered thoroughly in the theoretical literature (Hennart, 1988; Osborn and Baughn, 1990; Parkhe, 1993; Williamson, 1991). Less attention has been devoted, however, to the competitive advantages that these alliances create, including the ability for an incumbent to position itself for future entry by a competitor. Clayton and Jorgensen (2005) and Chen and Ross (2000), for instance, both provide theoretical models in which strategic alliances can be used by incumbents to discourage entry. Yet empirical studies have mainly focused on small or new firms who use alliances to establish an increased standing in the marketplace, rather than incumbents who seek to preserve their market share (Eisenhardt and Schoonhoven, 1996; Mitchell and Singh, 1992; Shan, 1990). To our knowledge, there has not been an attempt to empirically detect such motives behind the creation of strategic alliances.

The airline industry has been especially fertile ground for strategic alliances since the late 1990s, during which time numerous carriers have entered into cooperative marketing agreements called "codeshares."[1] Codesharing is a practice between carriers in an alliance which allows each carrier to sell seats on specified flights that are operated by its partners. Due to its increasing presence in the airline industry, codesharing has drawn the attention of anti-trust authorities who are concerned about possible anti-competitive motives.

From an econometrician's perspective, codesharing is a useful case to study because of the way in which the alliances between major U.S. "legacy" carriers have introduced codesharing on different flight segments at different times. This contrasts with a typical strategic alliance where two companies conduct a one-time merging of resources or efforts in a single market. Because the airlines have effectively formed their alliances multiple times, the resulting variation across markets and time provides the econometrician a unique opportunity to study which factors make markets conducive to alliance formation.

Like the strategic alliance literature in general, the literature on codesharing has explained the phenomenon in terms of two categories of incentives: (1) efficiency motives and (2) competitive motives. From the efficiency standpoint, codesharing effectively combines the networks of the allied airlines in order to provide an expanded availability and

☆ We would like to thank Jan Brueckner, Chris Foote, Dave Rapson, Ethan Singer, and seminar participants at the International Industrial Organization Conference for helpful comments and suggestions. The views expressed here are solely those of the authors and do not reflect the position of the Federal Reserve Bank of San Francisco, the Federal Reserve System, the U.S. Census Bureau, or the U.S. Department of Commerce.

* Corresponding author.
E-mail addresses: goetz@econ.umd.edu (C.F. Goetz), adam.shapiro@sf.frb.org (A.H. Shapiro).

[1] The term "codeshare" is derived from the fact that the airlines effectively share the flight by assigning each their own identification code, which allows them to list the flight separately in their respective reservation systems.

0167-7187/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ijindorg.2012.08.003

streamlining of service, a feature that has been written about in the airline networks literature (see, e.g. Brueckner and Spiller, 1991; Park et al., 2001). These efficiencies are exemplified in the "traditional" codeshared itinerary, in which a passenger flies on one partner's flight segment in order to connect to the codeshared flight segment that is operated by the other partner. Another efficiency gain from this arrangement stems from the cooperation between airlines removing the elevated markup, or "double marginalization", that is present when each airline sells its own flight segment separately. Since these network effects lead to lower cost and higher consumer welfare, social welfare is unambiguously increased.

The set of competitive motives for codesharing can be viewed as an attempt to gain or preserve market power, rather than to necessarily increase total surplus. One such theory put forth by Bamberger et al. (2004) suggests that codesharing gives allied airlines a marketing advantage over competitors by increasing the number of flight listings in computer databases, expanding the incumbent's passenger base, as well as sharing the marketing cost of attracting additional passengers. More generally, codesharing may be used as a commitment mechanism as demonstrated in the models of Spence (1977) and later Dixit (1980), who show how an incumbent firm can credibly reduce its price today when confronted with a potential entrant. In Dixit, the limit price is created by a leader firm's committed investment in capacity. However, the application to strategic alliances is perhaps better seen through Daughety's (1990) insight, who showed how a similar Stackelberg leadership position can be achieved by merging with another firm. In this respect, codesharing may be acting as a pseudo-merger which expands the incumbent's customer base and lowers future marginal cost.

In this study, we empirically test whether incumbent airlines are competitively motivated to form these strategic alliances in anticipation of potential entry by low-cost carriers (LCC). Using a within-flight segment specification of an airline's codesharing decision over time, we measure how the likelihood of an incumbent codesharing on a particular segment is affected when an LCC threatens to begin operating its own nonstop flight on the segment. Our method of identification rests on the notion of an exogenous threat of entry as defined by Goolsbee and Syverson (2008), where a competing LCC operates out of both endpoint airports of a flight segment at a particular point in time, but does not currently operate nonstop on the segment itself.

As this variable is a strong predictor of whether the potential low-cost entrant will begin operating their own nonstop flight on the segment in the future, we can identify competitive motives by measuring how alliance carriers react to its variations. This *preemptive competition* effect occurs when a threat of entry from a low-cost carrier induces the allied airline to begin codesharing. Specifically, because the threatening airline is likely to begin operating a nonstop flight on the segment in the near future, the incumbent operating airline may use codesharing as a tool to solidify their competitive standing in the market before entry occurs. Thus, this competition effect creates a positive relationship between a threat from outside the alliance and the probability of a codeshare. This identification strategy is based on the fact that there is little reason to expect efficiency motives to be any stronger or weaker in the face of such a threat. In other words, if consumer welfare gains and cost synergies are to be had by engaging in codesharing, then the allies should be codesharing regardless of the presence or absence of threats from third-party competitors, which do not affect these payoffs. Therefore, any preemptive action in response to the threat must be due to competitive motives.

Note that while such preemptive actions in the face of potential entry are consistent with entry-deterrence motives, our analysis will not be able to rule out the possibility that incumbents who form codeshares may alternatively be preparing to accommodate entry. While cost and efficiency factors do not change in the face of an exogenous threat, an incumbent might simply recognize that the price is going to be lower in the future when the LCC enters, and that the best way to accommodate the imminent entry is to begin codesharing now. This could give them the time needed to figure out how to best integrate the expanded customer base, (through changing flight schedules and gates, for instance), thereby making them better positioned to compete once actual entry occurs. Thus, we will consider both entry-deterrence and accommodation to be subsets of the preemptive competition motive.

Using a within-segment, fixed-effects linear probability estimator, we test for the presence of this preemptive motive. Indeed, we find that allied carriers have an increased propensity for codesharing when the segment is threatened by a low-cost carrier, providing evidence for the preemptive competition effect. Our estimates show that when a carrier is threatened on a segment by a low-cost carrier, the likelihood of codesharing that segment increases, on average, by a factor of about 25%.

To sharpen our identification, we test how the threat variables used in the previous specification interact with proxies for the degree of market power that the incumbent airline has. Specifically, by dividing the sample either according to market share levels or by whether the flight operates out of a hub, we can more clearly identify the motives behind the supply of codesharing. From a competitive standpoint, a larger market share or the presence of a hub means that the operator has "more to lose" if entry by the low-cost competitor does occur. Thus, if market-share preservation motives are in fact at work, then the rise in the probability of codesharing upon a threat should be more pronounced when the operator has higher market power. Our results indicate that both higher market share and the presence of a hub are strong determinants of the strength of the carrier's response to the threat from a low-cost competitor.[2] In all, our evidence perhaps sheds new light on the findings of Ito and Lee (2007), which show that the price of certain codeshared itineraries is generally lower than that of comparable non-codeshared tickets. We suggest that this may be due to the incumbent's use of a strategic alliance as a signal to potential entrants of their commitment to lower future prices, rather than due to a pure efficiency motive.

The remaining portion of the study is organized as follows. In Section 2, we give an overview of the rise of strategic alliances in the airline industry and the carriers' incentives for forming them. In Section 3 we review the data and in Section 4 we discuss our identification strategy. We explain our estimation method in Section 5 and discuss the results in Section 6. In Section 7 we perform robustness analysis and in Section 8 we perform an event analysis. We conclude in Section 8.

2. Overview of codesharing

2.1. The rise of codesharing

A codesharing agreement on a particular flight segment constitutes a strategic alliance in which the airline who operates the flight allows the partner airline to also sell seats on their flight. A "flight segment" is defined as a nonstop flight between two airports, and the carrier that provides the aircraft and crew for this flight is referred to as the "operating carrier" (or "operator"). Under a codesharing arrangement, the operator's flight will have multiple airlines selling its seats, and the carrier who sells a given passenger his ticket is called the "marketing carrier" (or "marketer"). Note that the codeshared segment can be used either as a passenger's sole flight to his destination, or as one leg in a multi-segment itinerary. Thus we call a "codeshared itinerary" any passenger's itinerary that contains at least one flight segment where the marketing carrier differs from the operating carrier.

For selling such itineraries, the marketing carrier of the segment receives a booking commission, while the operating carrier collects all the remaining proceeds from the fare. The exact arrangements vary; partners typically either authorize one another to sell a limited block of seats on a flight segment, or give the right to sell as many as the

[2] We note that the stronger response to a threat at a hub may also be indicating that airlines are using codesharing to bolster the competitive advantage that is already present at the hub. See Hendricks et al. (1997) and Aguirregabiria and Ho (2009).

other desires (Holloway, 2008). Additional costs of these arrangements include sharing and consolidating airport facilities (i.e. gates, lounges, maintenance equipment) and frequent flyer mile programs. Under U.S. laws, domestic carriers do not need to seek regulatory approval in order to codeshare on a given flight segment.[3]

Codesharing was ostensibly developed to increase network efficiency between partner carriers and became widespread in the international market during the 1990s as a way to offer connecting service for passengers in cities with limited access to international flights. For example, consider the longstanding codeshare between Air France and Delta on Air France's flight from Washington D.C. to Paris. This has enabled a passenger in Cincinnati to take a Delta plane to D.C. and connect to the Air France flight to Paris, all on a single itinerary purchased from Delta.[4] This setup effectively expands the route network for both airlines, and allows them to coordinate flight schedules and share certain airport facilities. It provides additional benefits to consumers who are spared the trouble of booking separate reservations on multiple airlines to get to their destination.[5]

As depicted in Fig. 1, codesharing expanded into the U.S. domestic market in the late nineties and exploded in 2003 with the formal partnering of United and US Airways in January. This was shortly followed by the alliance between Continental, Delta and Northwest in June of the same year. Airlines in these alliances became the exclusive codesharing partners for one another, although the decisions as to which flights would actually be codeshared were made on a case by case basis. The percentage of segments flown by legacy carriers that were codeshared reached a peak of around 34% in 2004. The purpose of domestic codeshared itineraries often mirrored that of the international ones. For instance, a Delta/Northwest codeshare on Delta's Atlanta–Minneapolis flight could allow Atlanta passengers to reach midwestern destinations on a single itinerary purchased from Northwest, by connecting from the codeshared Atlanta–Minneapolis flight to various Northwest flights originating in Minneapolis.

As the practice grew over time, with the number of passengers traveling on codeshare itineraries reaching around 2 million in 2005, the traditional codeshare was joined by other forms of codeshared itineraries. A study by Ito and Lee (2007) documented the importance of so-called "virtual" codeshares, in which the marketing airline of a passenger's itinerary is never the operating carrier for any segment of the trip. For example, a virtual codeshare occurs when a passenger buys a ticket for a United-operated flight from Philadelphia to Chicago via the US Airways website. While US Airways sells the ticket in this case, they are at no time involved in the passenger's actual flight. Ito and Lee found that such itineraries, which can be of both the nonstop and one-stop variety, represented over half of all codeshared itineraries by 2005 and were also cheaper than comparable non-codeshared itineraries. These facts suggest that the traditional rationale may not be the only reason why airlines engage in codesharing, therefore it is important to look at the different incentives that airlines face in deciding whether to codeshare a certain flight segment.

2.2. Incentives for codesharing

There are many reasons why an airline might choose to codeshare with its ally. For simplicity, we cast these reasons into two groups which represent the essential incentives behind strategic alliances in general: (1) efficiency motives and (2) competitive motives. The dichotomy here represents the fact that efficiency effects will always increase consumer surplus while competitive effects, depending on the outcome, can either raise or lower consumer surplus. We explain in more detail below.

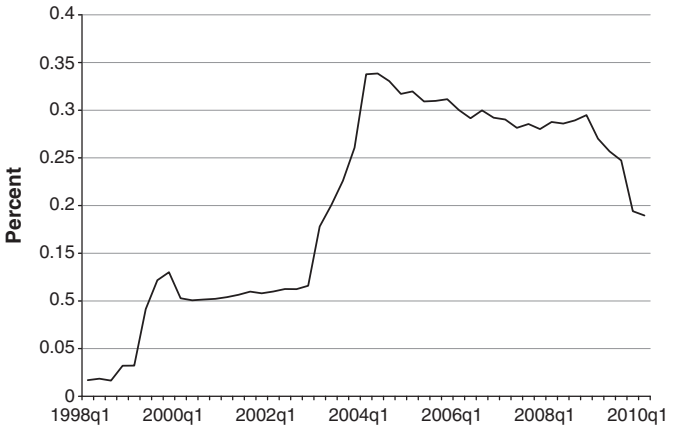

Fig. 1. Percent of all segments codeshared over sample period.

2.2.1. Efficiency motives

The first set of incentives stem from efficiency considerations, in that the codeshare improves overall welfare by enhancing the passenger's flying experience and/or lowering the airline's cost of operating the flight. The most obvious incentives are related to the traditional, one-stop codeshare structure which allows the airlines to integrate their networks and offer their passengers a wider range of destinations and schedules. This coordination of flights also increases welfare for the consumer by eliminating the need to buy two separate tickets, providing shorter layovers and gate transfers, as well as by streamlining check-in and baggage services.

Another important efficiency gain from traditional, one-stop codesharing occurs in the case where each partner airline holds some monopoly power in their respective markets. This efficiency gain stems from the fact that cooperation between airlines removes the "double marginalization" that is present when each airline sells their segments of the flight separately. As originally shown by Lerner (1934), when two tied goods are each sold by a different monopolist, the markup and deadweight loss are higher than when a single monopolist sells both of the goods. Therefore, when the two airlines work together to sell both segments of the flight, the price to the consumer should be lower than if the two airlines continued to control and market their own segments. Indeed, Ito and Lee (2007) found that traditional codeshare itineraries were cheaper than comparable non-codeshared ones.

While the efficiency gains from codesharing in the traditional inter-airline context are obvious, the motives for the virtual codesharing itineraries are less clear. Ito and Lee (2007) offer one explanation for virtual codesharing that relies on an efficiency rationale. They posit that airlines are using codeshares to price discriminate between passengers who prefer to buy from the operating carrier and those who are indifferent. For example, members of United's loyalty programs will be inclined to buy United flights directly through United in order to receive all the benefits of their membership, while price-sensitive consumers without such a brand loyalty will be willing to buy the more "generic" product from the allied airline.

The costs of codesharing are also an important consideration in determining whether to form an alliance on a certain flight segment. Codesharing arrangements often involve the consolidation and coordination of airport facilities, such as the rearranging of gates to facilitate passenger transfers, or the sharing of hangars and maintenance equipment. Additionally, arrival and departure times may need to be shuffled in order to achieve better flight schedule coordination. While these costs may not appear to be prohibitive they are clearly non-negligible,

[3] However, under 49 USC 41720, major carriers do have to give the Department of Transportation advance notice of their codeshares if they affect a large portion of their overall system capacity. We thank Peter Irvine of the DOT for this clarification.

[4] An itinerary of this type, where one marketing airline sells the entire itinerary, but each codesharing partner operates one flight segment of the trip, is dubbed a "traditional" codeshare by Ito and Lee (2007).

[5] These issues in international codesharing are explored in Brueckner (2001, 2003), Brueckner and Whalen (2000), and Park and Zhang (2000).

evidenced by that fact that less than 40% of alliance carrier's flights are deemed worthwhile to be codeshared.

2.2.2. Competitive motives

The second source of incentives for codesharing is related to an airline's desire to compete with other airlines. Like price and quality, codesharing is a tool that the airline can use to gain or protect market share from its competitors. For instance, Bamberger et al. (2004) suggest that codesharing gives the allied airlines a marketing advantage over their competitors. As a codeshared flight segment has two codes, it is listed twice on the electronic flight listings known as global distribution systems, which power most online marketing and reservations services like Expedia and Kayak. This means that a codeshared flight listing occupies more space on computer screens and can potentially crowd out competitors' listings. If both of the codesharing allies actually operate on the segment, this also leads customers to observe increased flight frequency, which may further boost the marketing benefit.

Codesharing may also give the alliance carriers a competitive advantage by increasing the proportion of brand loyal passengers traveling on the segment. Since the marketing carrier distributes frequent flyer miles to passengers traveling on the segment, there will exist passengers with frequent flyer travelers from both the operating carrier as well as the marketing carrier. This may act to increase the proportion of customers with large switching costs, thereby making it more difficult for the potential entrant to attract passengers.

Competitive reasons for codesharing may also be related to the seminal entry-deterrence model of Dixit (1980) which shows that if a potential competitor is threatening to enter the market, it can be in the incumbent's interest to preemptively expand its capacity, acting to reduce its future marginal cost of increasing output. In Dixit (1980) this is done with some capacity investment, however Daughety (1990) shows that the firm can achieve a capacity leadership position by merging with another firm. Increasing capacity either through investment or merger essentially constitutes a credible threat to lower price should the competitor choose to enter. Applied to our context, the pseudo merger with the codesharing partner may increase a carrier's capacity in the sense that it expands the number of potential passengers in the carrier's network. This is done by integrating the alliance partner's flights into the operating carrier's network. These new passengers may be both "virtual" codeshare passengers who fly the segment as a nonstop itinerary purchased from the partner carrier, as well as the "feeder" passengers who take the flight as one segment on a traditional codeshared itinerary. Expanding their customer base in this way may allow the incumbent airline to commit to lower prices, since marginal cost will likely fall due to economies of increased passenger density (Brueckner and Spiller (1991, 1994) and Berry et al. (2006)) and more efficient marketing.[6]

The above-mentioned sunk costs of codesharing are critical to the Dixit/Daughety story, since they serve as a credible signal of commitment to codesharing the flight. In other words, once the costs of consolidating airport facilities and flight schedules have been paid, the allied airlines have little reason not to continue codesharing in the future. Moreover, the fact that these costs are significant but not too large is also important, since Dixit explicitly notes that if costs exceed a certain threshold amount then it will not be in the incumbent's interest to make the capacity investment.

Note that these competitive motives could also act in tandem with the use of a hub as a base of market power as shown in Hendricks et al. (1997). In this model, the network complementarities originating from the presence of a hub allow the carrier to credibly lower prices on an individual segment without necessarily losing profits on the sale of one-stop itineraries. By increasing the presence at a hub, codesharing expands the networking capabilities of the incumbent carrier, thus bolstering this competitive effect.

Codesharing on a particular flight segment could also be used by alliance partners to contemporaneously compete against rivals in a variety of one-stop markets. In this case, it represents a joint attempt on behalf of the alliance carriers to divert market share away from a third party competitor's indirect service to one of the endpoint airports of the codeshared segment. This could obviously benefit the consumer by offering her more choices, possibly at a lower price. However, the tactic could also act to drive the third-party from the market, potentially lowering surplus. In our analysis, however, we will not be able to measure entry-deterrence on such one-stop route markets. Because of the way that our threat variable is defined by an airline having operations out of both endpoint airports, a threat and indirect service are often one and the same. That is, once an airline has established itself at the second endpoint airport, some indirect service is usually observed immediately. This leaves little opportunity to define an exogenous threat in the case of one-stop route markets.

Other competitive motives include more explicitly collusive goals as discussed in Gayle (2007, 2008) whereby codesharing could be an attempt to fix prices on overlapping routes. Nevertheless, collusion could also take a more subtle form as suggested by Bamberger et al. (2004). They hypothesize that codesharing could serve as a tacit agreement between partners not to expand into each other's markets. Allocating the codeshared segments evenly amongst each of the airline's flights acts as a commitment mechanism so that each party has equal incentive not to expand. Chen and Ross (2000) describe a similar scenario whereby the incumbent can profitably bribe a potential entrant into a more limited entry by sharing its facility at a low price.

3. Data and definitions

The goal of our analysis is to empirically test determinants of alliance carriers' codesharing decision. We do this by focusing on the effects of competitive *threats* on given flight segments, so it is necessary to measure both codesharing and competitive threats with our data. The data we use are from the Bureau of Transportation Statistics DB1B Origin and Destination survey, which collects information on a 10% sample of all airline ticket itineraries sold in the United States. Variables include the operating and marketing carriers for each flight segment of the trip, departure and arrival airports, stopover airports, number of passengers, fare paid, and other flight segment level characteristics. This is the same data used by Ito and Lee (2007), but we will extend their cross-sectional sample to a time-series one, spanning from 1998:Q1 to 2010:Q1. We collapse the data from the itinerary level to the quarterly segment–carrier level, so that each observation represents an operating carrier *i*'s nonstop flight segment *j* between two airports in time period *t*.[7]

In our analysis, the outcome variable of interest is whether a given operating carrier's flight segment between two airports is offered as a codeshare. Note that the codeshared passengers may either be using the flight as a self-contained nonstop trip (i.e. virtual codeshare passengers), or as a single leg in a multi-segment itinerary (i.e. traditional codeshare passengers). In our analysis, we will not distinguish between whether the flight segment is being sold predominantly as part of virtual codeshared itineraries or traditional ones, as almost all codeshared segments are sold as both types. How that codeshared segment is purchased, either as a virtual nonstop codeshare, a traditional one-stop codeshare, or some more complicated itinerary, is ultimately up to the consumer. Clearly, the airlines base their codesharing decisions on some knowledge of what kind of itinerary their customers are going to demand, but it is not evident from the data how to determine which passengers the codeshared segment is meant to facilitate.

[6] In the context of a segment, airlines must consider the cost of an additional flight or increasing the size of the aircraft. In this sense, increased passenger density may make it more cost effective to use larger aircrafts, which subsequently have lower marginal costs.

[7] For a more in-depth description of the construction of the data please see Appendix A.

While recognizing that the motives may differ depending on whether the codeshared flight is intended more for traditional or virtual itineraries, we believe that our specification models the codesharing decision in the most simple and direct way.

Therefore, for operating carrier i, flight segment j and time t, we define a binary codeshare dummy y_{ijt} indicating when at least one itinerary is sold on the flight segment where the marketing carrier is an alliance partner of the operating carrier, but not the operating carrier itself. This indicates to us that the airlines have a codesharing agreement on this flight segment.[8] We restrict the definition of codeshared flights to include only those by the major alliances partnerships listed in Ito and Lee (2007): Continental/Delta/Northwest, United/US Airways, American/Alaskan, American/Hawaiian, Northwest/Alaska, Continental/Alaska, Northwest/Hawaiian, and Continental/Hawaiian.

In order to quantify whether an airline's segment j is threatened by a potential low-cost entrant at time t, we follow the strategy of Goolsbee and Syverson (2008) and define an indicator variable for $LCCthreat_{jt}$ as follows: $LCCthreat_{jt}=1$ if *any* given low-cost airline (LCC) is operating flights out of both endpoint airports of segment j during time t, but is *not* actually operating a nonstop flight on segment j. If each given low-cost competitor is either flying on segment j, or is not operating at all out of at least one of the endpoint airports, then $LCCthreat_{jt}=0$. Goolsbee and Syverson use this as a proxy for threat of entry (for Southwest Airlines specifically),because when a low-cost airline has gates, desk agents, grounds crew and maintenance facilities established at both endpoint airports of a flight segment, it is easier to begin nonstop service on the segment. Indeed, the authors found that when Southwest Airlines threatened a market according to this definition, they were 18.5% more likely to enter the segment with a nonstop flight in the next quarter. In our analysis we extend the threat concept to also include the other low-cost airlines, Airtran, JetBlue, and Spirit.[9] Note that this threat-of-entry proxy is only appropriate for LCCs, due to the way in which these airlines are willing to fly routes between two non-hub airports. For traditional hub-and-spoke legacy competitors, however, the mere presence of operations in two airports is not a meaningful predictor of future nonstop service between the endpoints, since hub considerations are far more critical for such carriers. We do not attempt to measure "threats" coming from alliance partners, because provision of service to either endpoint by the codesharing partner is inherently endogenous to the partners' codesharing decision.[10] Furthermore, as noted above, we only measure threats on non-stop segments, as opposed to one-stop routes. This is because a threat of entry and the introduction of one-stop service are often one and the same. This leaves little opportunity to define an exogenous threat in the case of one-stop route markets.

Besides the threat proxies, we construct other variables on the carrier–segment–time level, such as the operating carrier's market share of non-codeshared passengers on the segment, and whether the endpoint airports are hubs for either the operating carrier or their allies. By interacting these variables with the threat proxy, we will be better able to see the competitive motive at work. After all the necessary variables are created, we restrict the sample to include only the members of the major codeshare alliances. Following Ito and Lee (2007) and Gerardi and Shapiro (2009), we ignore airline flight segments with low passenger counts in the DB1B sample, which we define as those with fewer than 10 observations in any given quarter or fewer than 100 over the entire sample period. Table 1 shows summary statistics from our final analysis data set.

4. Identification

Because we conduct a within-segment analysis, we exploit variation in codesharing status and threats on a given operator and segment over time. Identification of alliance airlines' motives for codesharing stems from the manner in which a threat will impact a carrier's decision to supply a codeshare on a particular segment at a given time. To aid in explaining this identification strategy, we explore the example depicted in Fig. 2. In this example, the codesharing decision concerns United's flight segment from Chicago to Omaha, which is depicted by the solid line connecting the two cities. Assume that the LCC Southwest currently has service out of Chicago with a flight to Raleigh, but as of now does not fly any flights out of Omaha. Therefore, United's codesharing decision should be based solely on efficiency concerns, as there is currently no competitive pressure.

To gain insight on the effects of a threat from a low-cost competitor, we examine what would happen if Southwest began flying to Omaha, thereby forming a "threat" to United on the Chicago–Omaha segment, as depicted by the dotted line in Fig. 2. Now that Southwest operates out of both Chicago and Omaha, they are more likely to enter the Chicago–Omaha segment with their own *nonstop* flight in the future. Note, however, that nothing has fundamentally changed on the Chicago–Omaha segment since Southwest is not yet an actual competitor.

Nevertheless, United may desire to codeshare the Chicago–Omaha segment with their alliance partner US Airways in order to prepare for Southwest's possible entry onto the Chicago–Omaha flight segment. The competitive effect of preemptively initiating this codeshare may stem from its positive impact on marketing, the increase in the proportion of brand loyal passengers, as well as the increase in traffic gained by attracting US Airways passengers connecting from Philadelphia. Therefore, if the probability of codesharing appears to increase once Southwest initiates a threat on the Chicago–Omaha segment, then we conclude that United's decision to codeshare is a direct competitive reaction to the increased probability of future competition on the Chicago–Omaha segment.

This identification stems from the fact that the incentive for cost reduction or improved service from a codeshare should not be affected one way or another by a new threat from a competitor. As the threat is generated from Southwest's St. Louis to Omaha segment (along with Southwest's Raleigh to Chicago segment), it is unlikely to be correlated with any supply-related factors that would make the Chicago–Omaha segment more efficient for United to codeshare. Thus, if the probability of codesharing the segment rises upon a threat from a low-cost carrier, then we can conclude that the competitive motive is at work.

It is also important to note that the preemptive competition effect should be more pronounced the more United is concerned about potential entry. This would be the case, for instance, if United was enjoying

Table 1
Flight level data: summary stats across all quarters.

	Mean	Std. dev.
Codeshare segment	.204	.403
LCC threat	.291	.454
Southwest threat	.083	.276
JetBlue threat	.060	.238
Spirit threat	.025	.157
AirTran threat	.147	.354
LCC entry	.157	.364
Southwest entry	.105	.306
JetBlue entry	.013	.115
Spirit entry	.013	.112
AirTran entry	.031	.173

[8] DOT confirms that the marketing carrier is always the airline that sold the passenger his ticket, and not simply one that may have rebooked a passenger or issued a rewards ticket on another airline

[9] We do not include Frontier and ATA airlines in our set of LCCs as these carriers act more like regional hub-and-spoke carriers who often codeshare with legacy partners. For instance, over the sample period, 50% of ATA's flights and 45% of Frontier's flights involve an interline or codeshare with another carrier. In contrast, .7% of Southwest's flights, .06% of Spirit's, 5% of AirTran's and none of JetBlue's flights involved a codeshare or interline flight.

[10] We will, however, define a threat variable for legacy competitors and allies as a falsification test in Section 7.

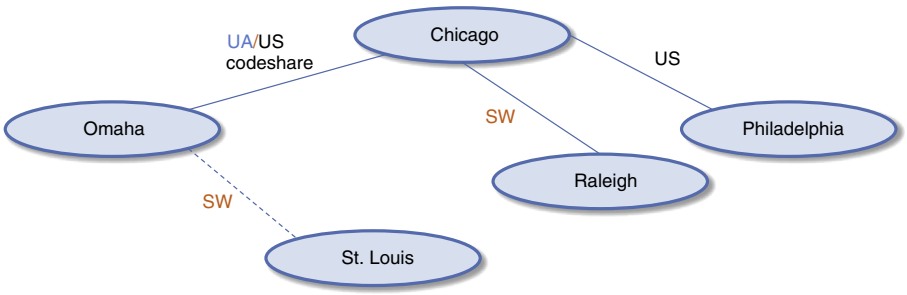

Fig. 2. Threats and codeshared flights.

high market share on the Chicago–Omaha segment. Generally, we expect the motive to be stronger when the operator has "more to lose" from increased competition. We explicitly test for the presence of this market-power effect in our estimation routine.

4.1. Identification concerns

Before delving into estimation, a discussion of other possible relationships between threats and the probability of codesharing is in order. There are three primary concerns. First, an initiation of a threat and an initiation of codesharing may both be responses to changes in aggregate demand for a certain destination. In this case, threats are not necessarily exogenous to the codesharing decision. Next, since an initiation of a threat means there are alternative ways of getting to a particular destination through indirect service, issues arise pertaining to how threats may affect the consumers' demand for codeshares. Finally, we discuss the possibility that the link between threats and codesharing is due to traditional competition over indirect service.

4.1.1. Aggregate demand effects

The first identification concern stems from our assumption that threats occur exogenously. In our example, the identification of the competition effect from Southwest's threat on United's codesharing comes from the fact that Southwest's decision to open service on the St. Louis–Omaha segment is exogenous to any codesharing considerations that United faces on the Chicago–Omaha segment. In other words, Southwest's decision to enter the St. Louis–Omaha segment must be uncorrelated with cost or demand factors in the Chicago–Omaha segment that would make it more suitable for United to codeshare. Since St. Louis–Omaha and Chicago–Omaha are clearly different markets, we see little reason why Southwest's decision to begin flying to Omaha (and therefore establish a "threat") would be directly motivated by changing factors in the Chicago–Omaha market. However, we recognize that endogeneity would arise if both the threat and codeshare decisions represent simultaneous responses to an increase in aggregate demand at an endpoint city (in this case, Omaha). We address this issue in our specification by controlling for the overall airport-level demand in the time period.[11]

4.1.2. Substitution effects

A threat may also have a more direct effect on consumer demand which leads to another identification concern. Namely, threats may affect demand through the codeshare's substitutability with alternative travel itineraries. For instance, consider an alternative scenario where the other half of Southwest's "threat" is represented by a flight from Philadelphia–Chicago (instead of Raleigh–Chicago), and where Southwest also happens to fly between Philadelphia and St. Louis, as depicted in Fig. 3. In this case Southwest's service conveniently serves Philadelphia–Omaha passengers through St. Louis, potentially reducing demand from passengers who might otherwise have purchased a codeshared Philadelphia–Chicago–Omaha itinerary from United/US Airways. While this convenient case might not typically occur, more generally, an LCC threat could represent new opportunities for passengers to reach the endpoint cities, which could result in less demand for codeshared segments. Thus, this "substitution effect" will induce a negative relationship between a threat and the probability of a codeshared itinerary being observed. Overall, this demand shifting will cause any estimate of relationship between threats and codeshares to be biased downwards. Since we are concerned with the probability of there being *any* codeshared itineraries on a flight, as opposed to the quantity of codeshares being purchased, this bias will be present when there are only a small number of codeshares being sold on a particular segment.

4.1.3. Indirect-competition effects

A related concern on the supply side also arises from the fact that both threats and codeshares represent the creation of new one-stop service to the endpoint airports. Specifically, we must consider the possibility that codeshares are created primarily to compete with low-cost carriers for market share on these kinds of routes. In this case a positive relationship between threats and codesharing may represent a more traditionally competitive reaction to existing competition, rather than an anticipatory reaction to threats of entry. For instance, as in Fig. 3, United and US Airways may codeshare the Chicago–Omaha segment to compete with Southwest's one-stop Philadelphia–St. Louis–Omaha service, by jointly offering their own one-stop service from Philadelphia to Omaha through Chicago.

While competing over indirect routes may very well be one motive for codesharing, we discount its importance for explaining a positive relationship between threats and codesharing in our regression specification. Clearly, threats and codeshares are likely to be positively correlated in the cross-section, as certain markets are more conducive to one-stop service than others. However, since our specification is identified off variations in threats on the same segment over time, this sort of competition would spur a positive correlation only if a competitor's new indirect service induces the allied airlines to offer a codeshared, indirect itinerary that they had no interest in providing before the competitor entered that market.

Furthermore, the fact that Southwest's threat conveniently forms service from Philadelphia to Omaha in this example – thus entering a route that UA/US might naturally want to cooperate on – is a unique case that is unlikely to be prevalent enough to bias the results. More realistically, the threat from Southwest might be formed with a new flight to Omaha from some other city, for instance Sacramento or Phoenix. In this case the formation of the threat would not create competition on any indirect routes to Omaha that UA/US would be interested in jointly serving with a codeshare.[12] Nevertheless, while the threat by Southwest may indeed create a new way for some passengers

[11] In the next section, we also conduct a robustness test using a proxy for expected demand.

[12] Note that back in the scenario depicted in Fig. 2, Southwest's threat does not create *any* new indirect routes.

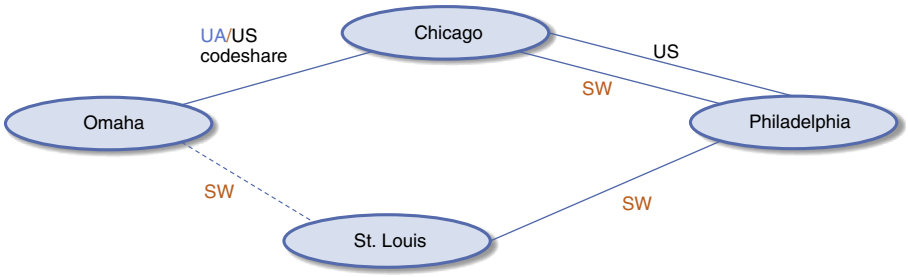

Fig. 3. Threats and indirect routes.

to get to Omaha via indirect service, it is unlikely that this option is a perfect substitute for United's service through Chicago.

In many cases, the one-stop flights formed by the threats and those formed by the codeshares are serving completely different sets of passengers. After all, the Chicago–Omaha codeshare can be used by passengers from many origin cities other than Philadelphia, as well as virtual nonstop passengers, a feature that is observed in the data. This makes it all the more implausible that a codeshare would be created primarily for the purpose of entering a specific one-stop route that a low-cost competitor is also entering. As mentioned earlier, increased demand for one-stop flights to Omaha might simultaneously lead to the formation of low-cost-carrier threats and codeshares, but the inclusion of overall airport-level demand in our regressions should alleviate this concern.

5. Estimation

As we are interested in estimating the marginal determinants of codesharing, we use a linear probability model (LPM). While index models, such as a probit and logit estimator, are advantageous in that they restrict the support of the probability space to the unit interval, they also introduce problems which can potentially lead to biased estimates of the marginal effects.[13] More generally, as discussed by Wooldridge (2002), the case for LPM is stronger when the independent variables are discrete, as in our case.[14]

5.1. Baseline fixed-effects specification

All of our specifications exploit the time-series variation in codesharing for a specific operator-segment, over and above the average probability of the operator codesharing on that segment and in that time period. Our baseline fixed-effects LPM specification is:

$$P\left(y_{ijt}=1\middle|X,\gamma\right)=\beta' X_{jkt}+\zeta' Z_{jkt}+\theta' Y_{jt}+\gamma_{ij}+\gamma_{it}+\varepsilon_{ijt} \quad (1)$$

where j indicates the segment, i indicates the operator, and y_{ijt} is a {0,1} variable indicating whether or not operator i is codesharing with an alliance partner on segment j in period t. Fixed, segment-specific factors such as distance and size of the endpoint cities are controlled for through operator-segment fixed intercepts represented by γ_{ij}. Time effects, specific to an operator, are controlled for with a set of operator-quarter dummies γ_{it}. We also control for the overall level of demand at the origin and destination airport with the vector Y_{jt}: two variables representing the total number of online passengers at the origin and destination airport, respectively.[15] The vector Zj_{kt} are variables indicating whether another carrier, k, actually operates on segment j. Our variables of interest lie in the vector X_{jkt}, a set of "threat" variables defined as in Section 3 — an occurrence where another carrier, k, operates out of both endpoints of the given segment, but does not operate a nonstop flight on the segment itself.[16] In all specifications, we cluster standard errors by segment.[17]

In our first estimation, our threat variable of interest represents a threat from any low-cost carrier (LCC):

$$P\left(y_{ijt}=1|X,\gamma\right)=\beta LCCThreat_{jt}+\zeta LCCEntry_{jt}+\theta' Y_{jt}+\gamma_{ij}+\gamma_{it}+\varepsilon_{ijt}. \quad (2)$$

In our second specification, we replace $LCCThreat_{jt}$ in Eq. (2) with the *number* of total threats, and in a third specification we create a separate dummy variable indicating a threat for each LCC in the sample. The preemptive competition effect will act to raise β as this indicates that the operator in question is increasingly likely to codeshare once a threat from a low-cost carrier occurs. However, as discussed above, the presence of the substitution effect will act to push β downwards for any of these threat variables. Thus, the sign of β will depend on relative sizes of the competitive and substitution effects.

6. Estimation results

6.1. Baseline specifications

We report results from our baseline fixed-effects specifications in Table 2. The first column shows the results of specification (2) and indicates that the coefficient on *LCCThreat* is .052 and significant at the 1% significance level. This result implies that codesharing increases by more than four and a half percentage points above the mean for a given operator's segment in a quarter when a low-cost carrier threatens the segment. Given a baseline probability of codesharing of 20% across the whole sample, this reveals that codesharing increases by roughly 25% on average when a segment is threatened. Also interesting is that the estimate on *LCCEntry* implies that the probability of codesharing stays elevated when the LCC actually enters the routes, implying that carriers stick with the codesharing agreement even though the precipitating reason was a threat of entry.

Our specification using the number of threats, in place of a threat dummy, produces similar results. The coefficient on the total number of LCC threats is .048. Finally, the estimates from decomposing the LCC threat variables into the individual operators are shown in the third

[13] The "fixed-effects probit" model, for instance, carries the "incidental parameters problem" in which case the parameter estimates are inconsistent unless the time frame is very large. One can partially correct for this using the fixed-effects logit which uses a conditional maximum likelihood estimator. However, we found this method to be computationally infeasible given the size of our data and the fact that we have two types of fixed effects.

[14] The LPM also has the advantage that its estimates will not be sensitive to the mean level of the predicted probability, since the marginal effect of each covariate is linear by construction.

[15] We also ran regressions where we included expected airport demand. Specifically, we included $\hat{Y}_{j,t+1}$ where the hat indicates we instrumented for this variable with current and a lagged value of overall airport demand. No results changed with this specification.

[16] See Appendix B for definitions of all of the variables used in estimation.

[17] Specifically, we cluster by "segment numeric" which avoids the double counting problem discussed in Gerardi and Shapiro (2009) of treating each direction as a separate segment. For example, BOS to LAX has the same segment numeric as LAX to BOS.

Table 2
Baseline LPM specification.

	(1)	(2)	(3)
LCCThreat	0.052***		
	(0.005)		
NumLCCThreat		0.048***	
		(0.004)	
SWThreat			0.050***
			(0.009)
JBThreat			0.051***
			(0.009)
ATThreat			0.051***
			(0.007)
SPThreat			0.049***
			(0.012)
LCCEntry	0.056***	0.060***	
	(0.007)	(0.007)	
SWEntry			0.067***
			(0.010)
JBEntry			0.003
			(0.017)
ATEntry			0.089***
			(0.014)
SPEntry			0.039**
			(0.017)
ln *OriginPass*	0.048***	0.048***	0.047***
	(0.006)	(0.006)	(0.006)
ln *DestPass*	0.041***	0.041***	0.041***
	(0.006)	(0.006)	(0.006)
Observations	344,293	344,293	344,293

Notes: all regressions include operator–segment and operator–time dummies. Standard errors are in parentheses and are clustered by segment-numeric to account for both autocorrelation and correlation between operators servicing the same two airports. One, two and three asterisks indicate significance at the 10%, 5% and 1% significance level, respectively.

column. All four LCCs have positive and significant coefficients and are similar in magnitude to the coefficient on the LCC threat variable. As before, these results can be interpreted as signs of a competitive motive, induced by threats from LCCs. To explore what drives these results, we now turn to our subsample specifications.

6.2. Subsample analysis

As mentioned above, we can gain more explicit identification of the supply-related motives of codesharing by exploring how market power affects the response of codesharing to the threat variables. Specifically, we expect the reaction to threats to be stronger on those segments where the operator has more at stake. To test this hypothesis, we perform two sets of subsample analysis. First we divide the sample based on the average marketshare of the operator on segment j over the course of the entire sample. Second, we divide the sample based on the hub status of the operator for the segment in question.

In the first subsample analysis, we divide the sample into three mutually exclusive bins: (1) those segments where the operator's average market share was below the 25th percentile (29.4%), (2) those segments where the operator's average market share was above the 75th percentile (98.4%), and (3) those segments which lied between the 25th and 75th percentile — the interquartile range (Table 3). If market-share preservation motives are present, then the effect of the threat from a competitor should be enhanced when the operator has a larger presence on the segment *j*. Results are reported in Table 3 and show that the within-segment effect of a threat on the probability of codesharing is largest on those segments in the upper 25th percentile of market share — essentially those segments where the operator is a monopolist. On segments where the operator's market power is low (those in the bottom 25th percentile), the effect of a threat of entry on the probability of codesharing is small and statistically insignificant.

To further test the importance of the market power on the response to a threat, we divide the sample into two mutually exclusive bins based on the hub status of the endpoint airports. Hubs represent airports where the operator has invested in creating a base of market power, and thus should serve as a good proxy.[18] Our results from the hub subsample analysis are shown in Table 4 and have similar implications to the subsamples defined based on average market share. Specifically, the positive coefficient on the threat variables is much more pronounced when at least one endpoint is a hub airport. In the sample of hub segments the effect of an LLC threat on the probability of codeshare is 0.076, compared to 0.028 in the non-hub sample. Overall, these results are indicative of the carrier responding to a stronger degree when the segment being threatened is servicing a hub airport.

Table 3
Subsamples based on average market share.

	Lower 25th	Interquartile range	Upper 25th
LCCThreat	0.010	0.061***	0.081***
	(0.006)	(0.007)	(0.013)
LCCEntry	0.021**	0.061***	0.101***
	(0.011)	(0.009)	(0.039)
ln *OriginPass*	−0.008	0.032***	0.073***
	(0.011)	(0.010)	(0.009)
ln *DestPass*	0.001	0.025***	0.058***
	(0.012)	(0.010)	(0.008)
Observations	86,082	172,149	86,062

Notes: subsamples were divided based on the average market share of the operator over the entire sample period. All regressions include operator–segment and operator–time dummies. Standard errors are in parentheses and are clustered by segment-numeric to account for both autocorrelation and correlation between operators servicing the same two airports. One, two and three asterisks indicate significance at the 10%, 5% and 1% significance level, respectively.

7. Robustness analysis

In this section we provide additional estimates from tests to evaluate the robustness of our findings. First, we perform a falsification test which is intended to test the validity of the threat measures. Second, we perform more general robustness tests by analyzing different subsamples as well as by including additional covariates and alternative measures of codesharing.

7.1. Falsification test

Our falsification test is based on the premise that Goolsbee and Syverson's (2008) threat measure pertains mainly to point-to-point operators, namely the low-cost carriers, as opposed to traditional hub-and-spoke carriers. This is because hub-and-spoke carriers are constrained by their hub-based networks. Thus, the act of a hub-and-spoke carrier operating out of both endpoints of a segment does not likely represent a true threat of future nonstop entry. For example, in Fig. 2, United will not find the threat of entry so credible if the threatening carrier were Delta instead of Southwest, since Delta does not have a hub in Chicago or Omaha. Generally, if a hub-and-spoke carrier does not fly on a specific segment, there is likely a network-related or historical reason behind it. This is why Goolsbee and Syverson (2008) chose to look only at threats from Southwest — an expanding low-cost carrier that operates a point-to-point network, a description that also applies to the LCCs included in our analysis.

Our second falsification test stems from the fact that a threat coming from a given carrier's alliance partner is inherently endogenous to competitive factors, and should therefore expose only efficiency motives or substitution effects. For example, in Fig. 2, if US Airways decided to begin a flight from Philadelphia to Omaha, then presumably US Airways would no longer need to send passengers to Omaha via Chicago through their codeshare agreement with United. Thus, if efficiency motives or

[18] See Borenstein (1989).

Table 4
Subsamples based on hub status.

	No hub	Hub
LCCThreat	0.028***	0.076***
	(0.005)	(0.007)
LCCEntry	0.030***	0.101***
	(0.006)	(0.010)
ln *OriginPass*	0.005	0.059***
	(0.005)	(0.008)
ln *DestPass*	0.008	0.050***
	(0.005)	(0.008)
Observations	136,789	207,504

Notes: all regressions include operator–segment and operator–time dummies. Standard errors are in parentheses and are clustered by segment-numeric to account for both autocorrelation and correlation between operators servicing the same two airports. One, two and three asterisks indicate significance at the 10%, 5% and 1% significance level, respectively.

Table 5
Falsification test.

	(1)	(2)
LCCThreat	0.052***	0.053***
	(0.005)	(0.005)
RivThreat	0.015*	0.017**
	(0.008)	(0.008)
AlliThreat		−0.049***
		(0.006)
LCCEntry	0.056***	0.061***
	(0.007)	(0.007)
RivEntry	−0.002	−0.003
	(0.004)	(0.004)
AlliEntry		0.092***
		(0.006)
ln *OriginPass*	0.048***	0.036***
	(0.006)	(0.006)
ln *DestPass*	0.041***	0.029***
	(0.006)	(0.006)
Observations	344,293	344,293

Notes: all regressions include operator–segment and operator–time dummies. Standard errors are in parentheses and are clustered by segment-numeric to account for both autocorrelation and correlation between operators servicing the same two airports. One, two and three asterisks indicate significance at the 10%, 5% and 1% significance level, respectively.

substitution effects are present, then a "threat" from a carrier's alliance partner should lower the probability that a given segment is codeshared.

Table 5 displays our baseline specification with the inclusion of a variable representing a threat from a rival legacy carrier – a legacy carrier outside of the observed carrier's alliance network – as well as a variable representing a threat from the observed carrier's alliance partner. The table shows that the threat from the rival has a negligible effect, and is statistically different from zero only at the 10% level, while the threat from an alliance partner is negative and statistically significant at the 1% level. However, the coefficient on the threat from the low-cost carrier remains at 0.05 and is statistically significant at the 1% level. The fact that the preemptive competition effect appears only as a result of threats by point-to-point carriers – precisely the type of carrier to which Goolsbee and Syverson (2008) meant for their threat variable to apply – gives further credence to the low-cost threat variable being a valid proxy for potential entry. In turn it is comforting that the "threats" from legacy competitors and allies, which don't reflect true entry threats, do not similarly generate positive coefficients.

7.2. Additional specifications and measures

In Table 6, we show six alternative regressions. We also include our baseline specification in the first column for ease of comparison. In the second and third columns, we add two measures of competition as independent variables to the baseline specification. Specifically, in the second column, we include a measure of the market concentration on the segment in the given time period, the Herfindahl–Hirschman Index (HHI), and in the third column we include the number of competitors operating on the segment in the given time period. These two measures serve as alternative indicators of the overall demand at each endpoint airport, and may therefore be correlated with the codesharing decision. In the fourth and fifth columns we provide a slightly different version of our codesharing dummy, y_{ijt}, which uses a higher threshold for deciding whether the segment is codeshared or not. Specifically, the fourth column shows estimates when we enforce that $y_{ijt}=1$ only if there are at least 5 codeshared tickets (representing 50 overall passengers) on the segment in the quarter, and the fifth shows results from setting $y_{ijt}=1$ only with a minimum of 10 such tickets. Finally, in the sixth and seventh columns we restrict our sample to the pre-2008 and pre-2005 time periods, respectively. Here we are addressing the fact that the Delta–Northwest merger in 2008 and the US Airways–America West merger in 2005 may have caused discrete and anomolous changes in the codesharing choices of these airlines. The resulting estimates indicate that the coefficients on the LCC threat variable are quite robust to all of these specifications. In particular, the estimate of *LCCThreat* continues to lie in the 0.05 range with very small standard errors.

8. Event analysis

As a final exercise we examine the timing of the codesharing response to the threat by looking at the periods before and after the initial threat occurs, as in Goolsbee and Syverson (2008). We look at the

Table 6
Robustness tests.

	Baseline	*HHI*	*NumComp*	Threshold five	Threshold ten	Pre 2008	Pre 2005
LCCThreat	0.052***	0.052***	0.054***	0.059***	0.061***	0.051***	0.042***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
LCCEntry	0.056***	0.046***	0.042***	0.055***	0.053***	0.050***	0.027***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.008)	(0.009)
ln *OriginPass*	0.048***	0.045***	0.041***	0.048***	0.047***	0.051***	0.045***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.007)	(0.006)
ln *DestPass*	0.041***	0.038***	0.033***	0.039***	0.035***	0.045***	0.035***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.007)	(0.006)
HHI		−0.069***					
		(0.011)					
NumComp			0.018***				
			(0.002)				
Observations	344,293	344,293	344,293	344,293	344,293	286,534	194,674

Notes: all regressions include operator–segment and operator–time dummies. Standard errors are in parentheses and are clustered by segment-numeric to account for both autocorrelation and correlation between operators servicing the same two airports. One, two and three asterisks indicate significance at the 10%, 5% and 1% significance level, respectively.

establishment of *initial* threats from the four LCCs in our sample in four separate regressions. The specification we use is:

$$P(y_{ijt}=1|X,\gamma)=\sum_{\tau=-8}^{8}\beta_\tau X^*_{jk,t_0+\tau}+\sum_{\tau=0}^{8}\zeta_\tau Z^*_{jk,t_0+\tau}+\theta' Y_{jt}+\gamma_{ij}+\gamma_{it}+\varepsilon_{ijt} \quad (3)$$

where $X_{jk,t_0+\tau}{}^*$ are time dummies surrounding the period when carrier k establishes a threat on segment j.[19] $Z_{jk,t_0+\tau}{}^*$ are time dummies at and subsequent to the period when carrier k establishes entry on segment j. Y_{jt} are two variables representing the total number of online passengers at the origin and destination airport, respectively. We limit our sample to segments where carrier k has had airport presence at either endpoint sometime during the sample period as well as segments where there exists at least one quarter of operation before the threat was established.[20] We again cluster our standard errors by segment and include both operator–quarter and carrier–segment fixed effects.

We report estimates of the four event study regressions in Table 7 as well as Fig. 4, which plots the point estimates on $X_{jk,t_0+\tau}{}^*$ along with 95% confidence intervals represented by the dotted lines. Our interest lies in comparing the deviation of the probability of codesharing from its trend probability (the omitted time variables which lie outside of the event window), before and after the threat is initiated on the segment. Note that since these regressions include operator-time fixed effects, any estimate that is not significantly different than zero should be interpreted as having no deviation from trend. If carriers are responding to a threat by codesharing, we would expect either the probability of codesharing to be below trend before the threat is initiated, or we would expected the probability of codesharing to be above trend around when the threat is initiated.

The estimates show that there is a rise in the probability of codesharing around the time period that the threat is established. For example, in the eight quarters before AirTran's initial threat, the probability of codeshare lies approximately on trend. However, in the period of the initial threat the point estimate rises to .019 and then .027 in the subsequent period. For Southwest, the rise in the probability of codesharing occurs in the period before the threat is established (0.052) and then rises in the subsequent three periods. This may be to due the fact that, as noted by Goolsbee and Syverson (2008), Southwest announces entry sometimes months before actual entry occurs. Overall, the estimates from the event analysis suggest that the probability of codesharing is larger after the threat is initiated than it was prior.

Table 7
Event analysis.

	Southwest	JetBlue	AirTran	Spirit
$Threat_{t-8}$	0.046***	−0.038***	0.006	−0.015
	(0.02)	(0.01)	(0.01)	(0.020)
$Threat_{t-7}$	0.021	−0.040***	0.01	−0.005
	(0.02)	(0.01)	(0.01)	(0.022)
$Threat_{t-6}$	0.022	−0.041***	−0.006	−0.022
	(0.01)	(0.01)	(0.01)	(0.025)
$Threat_{t-5}$	0.022	−0.029**	−0.003	0.004
	(0.02)	(0.01)	(0.01)	(0.024)
$Threat_{t-4}$	0.016	−0.036***	−0.006	0.022
	(0.02)	(0.01)	(0.01)	(0.025)
$Threat_{t-3}$	0.028*	−0.020*	−0.006	0.029
	(0.02)	(0.01)	(0.01)	(0.023)
$Threat_{t-2}$	0.018	−0.007	−0.001	0.045**
	(0.02)	(0.01)	(0.01)	(0.022)
$Threat_{t-1}$	0.050***	−0.007	0.007	0.059***
	(0.02)	(0.01)	(0.01)	(0.022)
$Threat_{t}$	0.053***	−0.01	0.019*	0.066***
	(0.02)	(0.01)	(0.01)	(0.022)
$Threat_{t+1}$	0.079***	0.013	0.027***	0.060***
	(0.02)	(0.01)	(0.01)	(0.022)
$Threat_{t+2}$	0.067***	0.001	0.018*	0.054**
	(0.02)	(0.01)	(0.01)	(0.022)
$Threat_{t+3}$	0.014	0.004	0.025***	0.081***
	(0.02)	(0.01)	(0.01)	(0.023)
$Threat_{t+4}$	0.008	−0.019	0.029***	0.058***
	(0.02)	(0.02)	(0.01)	(0.022)
$Threat_{t+5}$	0.031	0.001	0.031***	0.039*
	(0.02)	(0.01)	(0.01)	(0.020)
$Threat_{t+6}$	0.007	0.004	0.044***	0.040**
	(0.02)	(0.01)	(0.01)	(0.020)
$Threat_{t+7}$	−0.017	0.033**	0.030***	0.025
	(0.02)	(0.01)	(0.01)	(0.019)
$Threat_{t+8}$	−0.015	0.01	0.033***	0.006
	(0.02)	(0.01)	(0.01)	(0.019)
$Entry_{t}$	0.017	−0.037	0.02	0.076**
	(0.02)	(0.04)	(0.03)	(0.030)
$Entry_{t+1}$	0.032	−0.076**	0.024	0.072**
	(0.03)	(0.03)	(0.03)	(0.033)
$Entry_{t+2}$	0.052*	−0.086**	0.031	0.067*
	(0.03)	(0.04)	(0.02)	(0.035)
$Entry_{t+3}$	0.024	−0.099**	0.033	0.075**
	(0.03)	(0.04)	(0.03)	(0.033)
$Entry_{t+4}$	0.007	−0.088**	0.038	0.060*
	(0.03)	(0.04)	(0.03)	(0.034)
$Entry_{t+5}$	0.019	−0.074	0.003	0.062*
	(0.03)	(0.05)	(0.03)	(0.032)
$Entry_{t+6}$	−0.029	−0.025	0.039	0.036
	(0.03)	(0.04)	(0.03)	(0.034)
$Entry_{t+7}$	−0.013	−0.048	0.02	0.040
	(0.03)	(0.04)	(0.03)	(0.036)
$Entry_{t+8}$	−0.018	−0.041	0.026	−0.025
	(0.03)	(0.06)	(0.03)	(0.033)
Observations	23,473	37,478	51,314	14,826

Notes: all regressions include operator–segment and operator–time dummies. Standard errors are in parentheses and are clustered by segment-numeric to account for both autocorrelation and correlation between operators servicing the same two airports. One, two and three asterisks indicate significance at the 10%, 5% and 1% significance level, respectively.

9. Conclusion

Our analysis provides evidence for strategic alliances being used by legacy airlines as a competitive tool to preemptively prepare for potential entry. The results from our fixed-effects linear probability estimates show that the probability of allied airlines codesharing on a given route increases in the face of threats from low-cost competitors. They also reveal that this effect is magnified when the incumbent airline has high market share or a hub at one of the endpoint airports, further indicating that airlines codeshare with the motive of preserving their market power against potential entrants.

These results have interesting implications for the results found by Ito and Lee (2007), which showed that virtual codeshare itineraries are priced lower, ceteris paribus, than online itineraries. Our estimates indicate that this lower price may be the outcome of codesharing being used as a method of competitive behavior, rather than, as the authors explain, the codeshared itinerary being a differentiated, lower-quality product. As the Dixit framework suggests, lower prices are the natural byproduct of the incumbent firm's pre-commitment to compete. By combining their marketing, scheduling, and airport facility capacity with an ally, they are better able to commit to a lower price in anticipation of future competition. This may, in turn, also help explain the Goolsbee–Syverson finding that incumbent airlines lower prices when a nonstop market is threatened by Southwest Airlines. By expanding frequent flyer programs for passengers who travel a certain segment, codesharing potentially makes these passengers more loyal by increasing their switching costs, thus making the associated price drop more credible.

In sum, the use of codesharing as a strategic alliance may indeed lead to network efficiencies and prices that are generally lower than

[19] As in Goolsbee and Syverson (2008), these post-t_0 dummies take a value of one only if the LCC has not yet entered the segment.

[20] Due to this restriction, the four samples (one for each of the four LCCs) analyzed here are smaller than the sample used in our baseline fixed-effects analysis.

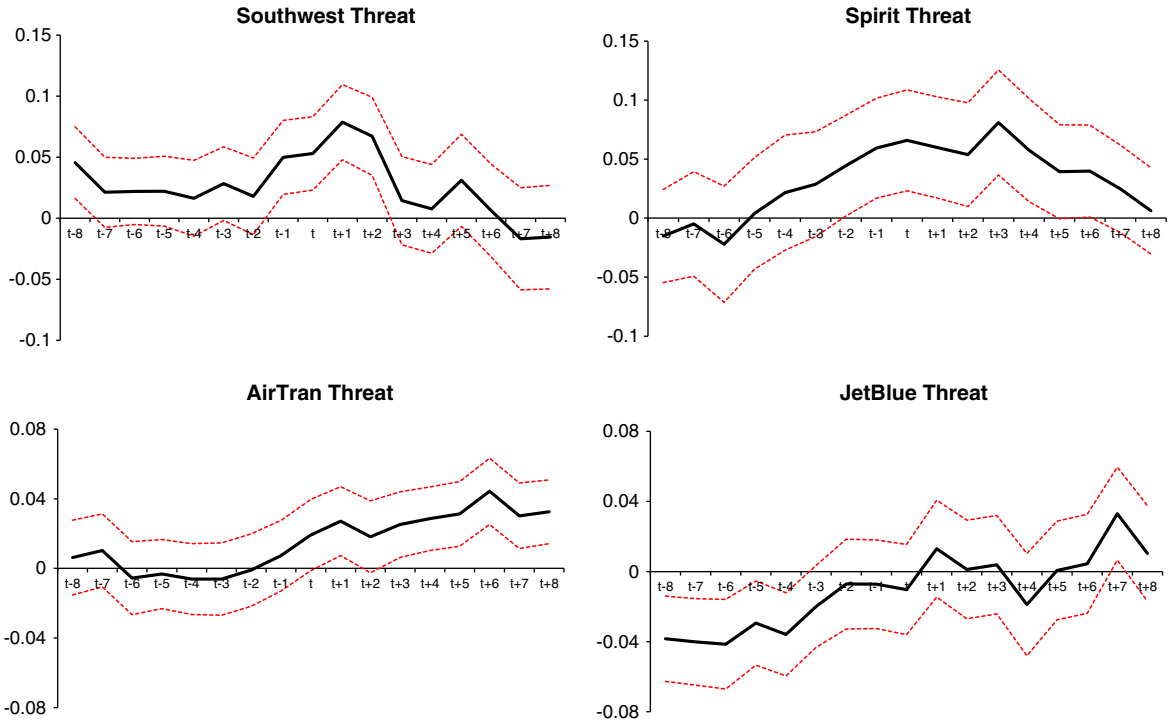

Fig. 4. Event analysis. All regressions include operator-time as well as operator-segment fixed effects. 95% confidence intervals are indicated by dotted lines and represent standard errors clustered by segment-numeric to account for both autocorrelation and correlation between operators servicing the same two airports.

non-codeshared flights, but it may also deter future competition on routes controlled by legacy airlines. We leave for future research a more detailed welfare analysis, which would be needed to determine whether codesharing has overall had a positive impact on consumer and social welfare. Our study demonstrates, however, the importance of considering the role of strategic alliances such as codeshares not simply in terms of a one-time combining of resources and networks, but rather as a dynamic response to the competitive landscape over time.

Appendix A. Data construction

Here we describe the steps used to construct our panel data set, which is built on the airline–segment–quarter level of observation. All raw data are from the Bureau of Transportation Statistics DB1B Origin and Destination survey, which represents a 10% sample of all domestic airline tickets and is available at the TranStats website. The DB1B data set is published quarterly as three separate files, representing coupon data, market data, and ticket data, which can be merged together by a unique itinerary id.[21] The raw data are constructed on the coupon-level, meaning there is a separate observation for each flight segment that a passenger flies between two airports. The key variables of interest for our purposes are the operating carrier ("opcarrier") and the marketing carrier ("tkcarrier") of each coupon that a passenger purchases.

We limit our sample to domestic coach-class flights from one-way and roundtrip itineraries. We filter out tickets whose price is of questionable magnitude, as they may represent frequent flyer reward tickets or other types of unique passengers, whose observed marketing carrier may not be appropriate for the purpose of detecting a codeshare. The DB1B includes the variable "dollar_cred", which flags itineraries with unreliable fares according to a set of criteria defined by BTS, and we delete itineraries that are flagged as non-credible. We also delete tickets flagged by the variable "bulk_fare", which are tickets packaged by a third party seller. In addition, we delete itineraries whose fare is less than $10 per directional trip, as these are likely to be award tickets from a frequent flyer program.

After filtering the itineraries, we collapse all flight coupons within a given quarter that share the same origin and destination, and same combination of operating and marketing carriers. For instance, say we observe 1000 coupons in 2000:q1 of passengers who flew on United from Chicago to Omaha, and purchased the flight from United. Then we would combine these into a single observation which records the number of such itineraries. Suppose we also observe another 50 coupons on the United Chicago–Omaha flight that are marketed by US Airways, in which case we would create a second observation documenting these cases.

Before we can accurately identify codeshared segments, we must first make a few modifications to account for some special cases. First of all, legacy carriers often engage small regional airlines to operate their flights, either under the name of the legacy carrier or their own. Even though the marketing carrier may differ from the operating carrier in the data, we do not want to count the flights as codeshared, since the regional airlines are really acting as subsidiary airlines and not as codesharing allies. To do this, we follow Ito and Lee (2007) recode the regional carriers to be the same as the legacy carrier that they are operating under. We make similar modifications for mergers, when the two airlines are officially one company, but the subsumed airline continues to operate under its own code for some time afterward. Specifically we recode America West to be US Airways starting in 2005:q3, and we recode Northwest to equal Delta starting in 2008:q4. Now we can properly determine whether a flight segment truly has marketing carriers that differ from the operating carrier.

[21] For further details on the data source, see the BTS's website http://www.transtats.bts.gov.

Finally, we collapse all the information in each quarter by operating carrier-flight segment. That is, a single data point is created for each operating carrier's flight, and includes variables recording all the marketing carriers for this flight. In our above example, the two observations we have for United Chicago–Omaha flight are combined into one, and information on the 1000 tickets sold by United and the 50 sold by US Airways is recorded. Now we can use this information to determine whetheran operating carrier's flight segment is codeshared with an alliance airline at each period in time. The alliance groups that we study are American/Alaska/Hawaii, Continental/Delta/Northwest, and United/US Airways.

This operator-segment level data is also used to calculate the threat variables, since we now have a database of all the nonstop flight segments that each airline operates. For instance, to calculate whether Southwest threatens United's Chicago–Omaha flight segment, we check to see whether Southwest operates any flights out of the Chicago and Omaha airports, as well as whether they operate nonstop on the Chicago–Omaha segment.

Once all the necessary variables are created, we limit the final regression sample to include only the major legacy carriers who engage in codesharing, since these are the airlines whose decision we are analyzing. That is, we only keep observations for the carriers in the major alliance groups listed above. We also eliminate routes with very low passenger counts, as these routes are very small market routes that are often served by only a single regional carrier. Such routes have virtually no chance of being threatened by LCCs or codeshared. Our thresholds require that an airline must have at 10 observed passengers in a quarter and at least 100 passengers over the entire sample period in order to be included.

Appendix B. Variable definitions

- $y(X,\gamma)_{ijt}$ — codesharing dummy variable which indicates if at least one ticket was sold on airline *i*'s nonstop flight segment *j* in quarter *t* by a marketing carrier other than *i*, but who is a member of *i*'s alliance network.
- X_{jkt} — set of threat dummy variables, which indicate whether a given airline (or one of a group of airlines), *k*, operates out of both endpoint airports of segment *j* at time *t*, but does not operate nonstop service on segment *j* itself. The specific threat variables are as follows:
  - $LCCThreat_{jt}$ — indicates whether nonstop segment *j* is threatened at time *t* by at least one of low cost carriers: Southwest, JetBlue, AirTran, or Spirit.
  - $SWThreat_{jt}$ — indicates whether segment *j* is threatened at time *t* by Southwest Airlines.
  - $JBThreat_{jt}$ — indicates whether segment *j* is threatened at time *t* by JetBlue.
  - $ATThreat_{jt}$ — indicates whether segment *j* is threatened at time *t* by Air Tran.
  - $SPThreat_{jt}$ — indicates whether segment *j* is threatened at time *t* by Spirit Airlines.
  - $RivThreat_{jt}$ — indicates whether nonstop segment *j* is threatened at time *t* by at least one legacy airline outside of its alliance network.
  - $AlliThreat_{jt}$ — indicates whether nonstop segment *j* is threatened at time *t* by at least one legacy airline inside of its alliance network.
  - $NumLCCThreat_{jt}$ — indicates the total number of threats on segment *j* at time *t* stemming from low-cost carriers.
- Z_{jkt} — Set of entry dummy variables, which indicate whether a given airline, *k*, operates out of both endpoint airports of segment *j* at time *t*. The specific entry variables are as follows:
  - $LCCEntry_{jt}$ — indicates whether nonstop segment *j* is being operated at time *t* by at least one of low cost carriers: Southwest, JetBlue, AirTran, or Spirit.
  - $SWEntry_{jt}$ — indicates whether segment *j* is being operated at time *t* by Southwest Airlines
  - $JBEntry_{jt}$ — indicates whether segment *j* is being operated at time *t* by JetBlue.
  - $ATEntry_{jt}$ — indicates whether segment *j* is being operated at time *t* by Air Tran.
  - $SPEntry_{jt}$ — indicates whether segment *j* is being operated at time *t* by Spirit Airlines.
  - $RivEntry_{jt}$ — indicates whether nonstop segment *j* is being operated at time *t* by at least one legacy airline outside of its alliance network.
  - $AlliEntry_{jt}$ — indicates whether nonstop segment *j* is being operated at time *t* by at least one legacy airline inside of its alliance network.
- ln $OriginPass_{jt}$ — the total number of non-codeshared passengers traveling through the origin airport of segment *j* in period *t*.
- ln $DestPass_{jt}$ — the total number of non-codeshared passengers traveling through the destination airport of segment *j* in period *t*.
- HHI_{jt} — the Herfindahl–Hirschman Index of segment *j* in period *t* calculating using passenger counts in the *DB1B* data.
- $NumComp_{jt}$ — the number of competitors on segment *j* in period *t* calculated using the *DB1B* data.

References

Aguirregabiria, V., Ho, C.-Y., 2009. A dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments. Manuscript. University of Toronto.

Bamberger, G., Carlton, D., Neumann, L., 2004. An empirical investigation of the competitive effects of domestic airline alliances. Journal of Law and Economics XLVII, 195–222.

Berry, S., Carnall, M., Spiller, P., 2006. Airline hubbing, costs and demand. Advances in Airline Economics, Competition Policy and Anti-Trust pages 1, 183–214.

Borenstein, S., 1989. Hubs and high fares: dominance and market power in the U.S. airline industry. The RAND Journal of Economics 20 (3), 344–365.

Brueckner, J., 2001. The economics of international codesharing: an analysis of airline alliances. International Journal of Industrial Organization 19, 1475–1498.

Brueckner, J., 2003. International airfares in the age of alliances: the effects to codesharing and antitrust immunity. The Review of Economics and Statistics 85, 105–118.

Brueckner, J., Spiller, P., 1991. Competition and mergers in airline networks. International Journal of Industrial Organization 9, 323–342.

Brueckner, J., Spiller, P., 1994. Fares and network 'feed': estimating economies of traffic density in airline hub-and-spoke networks. Journal of Law and Economics.

Brueckner, J., Whalen, T., 2000. The price effects of international airline alliances. Jounal of Law and Economics 43, 503–545.

Chen, Z., Ross, T., 2000. Strategic alliances, shared facilities, and entry deterrence. The RAND Journal of Economics 31 (2), 326–344.

Clayton, M., Jorgensen, B., 2005. Optimal cross holding with externalities and strategic interactions. Journal of Business 78, 1505–1522.

Daughety, A., 1990. Beneficial concentration. American Economic Review 80 (5), 1231–1237 (December).

Dixit, A., 1980. The role of investment in entry-deterrence. The Economic Journal 90 (357), 95–106 (March).

Eisenhardt, K., Schoonhoven, C., 1996. Resource-based view of strategic alliance formation: strategic and social effects in entrepreneurial firms. Organization Science 7 (2), 136–150.

Gayle, P., 2007. Airline code-share alliances and their competitive effects. Journal of Law and Economics 50, 781–819.

Gayle, P., 2008. An empirical analysis of the competitive effects of the Delta/Continental/Northwest code-share alliance. Journal of Law and Economics 51 (4), 743–766 (November).

Gerardi, K., Shapiro, A.H., 2009. Does competition reduce price dispersion? New evidence from the airline industry. Journal of Political Economy 117 (1), 1–37.

Goolsbee, A., Syverson, C., 2008. How do incumbents respond to the threat of entry? Evidence from the major airlines. Quarterly Journal of Economics 123 (4), 1611–1633 (November).

Hendricks, K.M., Piccione, Tan, G., 1997. Entry and exit in hub-spoke networks. The RAND Journal of Economics 28, 291–303.

Hennart, J., 1988. A transaction cost theory of equity joint ventures. Strategic Managment Journal 123 (4), 1611–1633 (November).

Holloway, S., 2008. Practical Airline Economics. Straight and level, third ed. Ashgate Publishing, Ltd.

Ito, H., Lee, D., 2007. Domestic code sharing, alliances, and airfares in the U.S. airline industry. Journal of Law and Economics 50, 355–380.

Lerner, A.P., 1934. The concept of monopoly and the measurement of monopoly power. Review of Economic Studies 1, 157–175.

Mitchell, W., Singh, K., 1992. Incumbents' use of pre-entry alliances before expansion into new technical subfields of an industry. Journal of Economic Behavior and Organizations 18, 347–372.

Osborn, R., Baughn, C., 1990. Forms of interorganizational governance for multinational alliances. Academy of Management Journal 33, 503–519.

Park, J.H., Zhang, A., 2000. An empirical analysis of global airline alliances: cases in the North Atlantic markets. Review of Industrial Organization 16, 367–384.

Park, J.H., Zhang, A., Zhang, Y., 2001. Analytical models of international alliances in the airline industry. Transportation Research Part B 35, 865–886.

Parkhe, A., 1993. Strategic alliance structuring: a game theoretic and transaction cost examination of interfirm cooperation. Academy of Management Journal 36, 794–829.
Shan, W., 1990. An empirical analysis of organizational strategies by entrepreneurial high-technology firms. Strategic Management Joumal 11, 129–139.
Spence, A.M., 1977. Entry, investment, and oligopolistic pricing. Bell Journal of Economics 8 (Fall), 534–544.
Williamson, O., 1991. Comparative economic organization: the analysis of discrete structural alternatives. Administrative Science Quarterly 36, 269–296.
Wooldridge, J., 2002. Econometric Analysis of Cross Section and Panel Data. MIT Press.