

American Economic Association

Adaptation and Vertical Integration in the Airline Industry

Author(s): Silke Januszewski Forbes and Mara Lederman

Reviewed work(s):

Source: *The American Economic Review*, Vol. 99, No. 5 (Dec., 2009), pp. 1831-1849

Published by: [American Economic Association](#)

Stable URL: <http://www.jstor.org/stable/25592538>

Accessed: 07/01/2013 17:16

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The American Economic Review*.

<http://www.jstor.org>

Adaptation and Vertical Integration in the Airline Industry

By SILKE JANUSZEWSKI FORBES AND MARA LEDERMAN*

We explore patterns of vertical integration in the US airline industry. Major airlines subcontract portions of their network to regional partners, which may or may not be owned. We investigate if ownership economizes on ex post renegotiation costs. We estimate whether airlines are more likely to use owned regionals on city pairs with adverse weather (which makes adaptation decisions more frequent) and on city pairs that are more integrated into the major's network (which raises the costs of having adaptation decisions resolved suboptimally). Our results suggest a robust empirical relationship between adaptation and vertical integration in this setting. (JEL L14, L22, L24, L93)

How do firms decide which transactions to carry out in-house and which to procure through the market? In response to this question, a large theoretical literature has developed. One branch of this literature—transaction cost economics—argues that integration may be a more efficient way to organize when contracts are incomplete and parties cannot easily switch trading partners (Oliver Williamson 1975, 1985).¹ Under these conditions, ex post renegotiation will be both necessary and costly and, as a result, the transaction costs of market exchange may be substantial.

Empirical studies testing this hypothesis have focused primarily on the effect of asset specificity on vertical integration and have generally found a positive relationship.² However, as Steven Tadelis (2002) emphasizes, holding asset specificity constant, variation in transaction costs will also result from differences in transaction complexity. Complex transactions are associated with greater contractual incompleteness, and in turn more costly haggling over ex post adaptation decisions. Yet, relative to asset specificity, the frequency of adaptation decisions as a determinant of vertical integration has received considerably less attention in the empirical literature.³

*Forbes: University of California, San Diego, 9500 Gilman Dr., 0508, La Jolla, CA 92093 (e-mail: sjanuszewski@ucsd.edu); Lederman: Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, Ontario, Canada M5S 3E6 (e-mail: mara.lederman@rotman.utoronto.ca). An earlier version of this paper was circulated under the title “Control Rights, Network Structure, and Vertical Integration: Evidence from Regional Airlines.” We thank two anonymous referees, Daniel Akerberg, Mike Arguello, Pierre Azoulay, George Baker, Severin Borenstein, Vincent Crawford, Ken Corts, Martin Hellwig, Barry Hirsch, Ig Horstmann, Thomas Hubbard, Cathy Jacobs, Darin Lee, Scott Masten, Scott McCartney, Wallace Mullin, Nancy Rose, Amedeo Odoni, Brian Silverman, Joel Sobel, Andrew Sweeting, Steven Tadelis, Dee Temples, and seminar participants at MIT, UCLA, UCSD, the UBC Sauder School of Business, the Rotman School of Management/Institute for Policy Analysis Conference on Organizational Economics, the 2005 IIOC, and the 2006 NBER IO Winter Meeting for helpful comments and conversations. All errors are our own.

¹ The other branches of this literature are the property rights theory of the firm (see Sanford Grossman and Oliver Hart 1986 and Hart and John Moore 1990) and agency theory (see Bengt Holmström and Paul Milgrom 1991 and Holmström 1999). See Robert Gibbons (2005) for a summary and attempt at integration.

² See, for example, Kirk Monteverde and David Teece (1982), Erin Anderson and David Schmittlein (1984), Scott Masten (1984), Paul Joskow (1985), and Thomas Hubbard (2001).

³ Two of the studies cited above also include measures of complexity. Masten measures an aerospace component's complexity with a three-level ranking scheme that was developed by the company itself. Monteverde and Teece measure an automobile component's complexity with the amount of engineering effort that was required to design that component. Both studies find that more complex transactions are more likely to be integrated.

This paper examines patterns of vertical integration in an industry in which transactions are complex, contracts are incomplete, and ex post adaptation occurs frequently. Our setting is the US regional airline industry. All of the large network airlines, often called “majors,” subcontract some of their flights to regional carriers. There is substantial heterogeneity both within and across majors in the extent to which these regional partners are owned. While majors and regionals write contracts that specify the schedule of flights that the regional is to operate, the complexity of airline operations prevents them from writing contracts that specify how that schedule would change as different contingencies arise. As a result, ex post adaptation decisions become necessary, resulting from factors such as adverse weather, mechanical problems, or air traffic control disruptions. However, majors and regionals will generally face conflicting incentives with respect to schedule adjustments. While majors seek to optimize the overall profitability of their network, regionals face financial incentives that are based only on the subset of routes that they serve. Thus, the costs of renegotiation with an independent regional may be large, possibly creating a role for ownership as a means to facilitate ex post adaptation.

To investigate this, we estimate whether airlines are more likely to use owned regionals to serve city pairs on which adaptation decisions need to be made more frequently and on city pairs on which it is more costly to have these decisions resolved in a less-than-optimal way. We exploit the fact that both of these vary with observable market characteristics. First, we use the average weather patterns at the endpoint airports of a city pair to measure the probability that flights on a particular city pair will be subject to ex post schedule adjustments. Second, we use the degree to which a given city pair is integrated into the major’s overall network to measure the costs to a major of having schedule adjustments resolved suboptimally. Flights on more integrated city pairs are more likely to impose externalities on other flights in the major’s network because they compete with these flights for access to scarce airline and airport facilities, as well as carry more connecting passengers.

Our empirical analysis shows that vertically integrated regionals are more likely to be used on routes between airports with more precipitation and, in particular, with more snowfall. In addition, we find that owned regional partners are more likely to be used on city pairs that have the major’s hub at least at one endpoint and on city pairs between airports at which the major operates a larger number of flights. One advantageous feature of our data is that several majors own some of their regional partners while also contracting with others. This allows us to include firm dummies and investigate the factors affecting the “make-or-buy” decision for a *given* firm. We find similar effects with and without airline dummies in the estimation, which suggests that our effects are identified by variation *within* firms as well as variation across firms.

This paper contributes to the literature that seeks to explain patterns of vertical integration within industries (see Francine Lafontaine and Margaret Slade 2007, for a review). Most significantly, we believe this study is the first to carry out a large-scale analysis of the relationship between integration and the costs of ex post adaptation using explicit measures of both the likelihood and costliness of adaptation decisions. The clear relationship between weather and flight disruptions provides us with a way to measure the probability of ex post adaptations in a precise and credible way. In addition, our setting allows us to investigate the relationship between vertical integration and underlying transaction characteristics, while controlling for unobservable factors that affect a particular firm’s relative returns from vertical integration. To our knowledge, this is one of the first papers outside of the franchising literature to examine this.

Finally, our paper shows how network considerations can affect integration decisions. The network structure of airlines’ operations is at least partially responsible for the complexity of transactions in our setting. Much of the reason why majors and regionals cannot write complete contracts is because the optimal response to any schedule disruption will have to be conditioned on events taking place elsewhere in the airline’s network. The set of possible contingencies is

so large that it would be prohibitively costly for the major and the regional to specify a contract covering all of them. In addition, the reason why haggling between majors and regionals over schedule changes may be costly is that while majors seek to optimize the overall profitability of their network, it is difficult for them to use contracts to provide regionals with incentives that extend beyond the small set of routes they serve.⁴ We suspect that both of these issues may arise when firms subcontract part of their operations in other network-based industries.

The remainder of this paper is organized as follows. The next section provides relevant institutional detail. In addition to explaining why and how the need for ex post adaptation arises, we describe the likely costs of ownership that result from idiosyncratic features of labor relations in this industry. Section II presents our empirical approach. In Section III we describe the data and variables, and in Section IV we present the results. A final section concludes.

I. Institutional Background

A. *The Role of Regional Airlines*⁵

Regional airlines operate as “subcontractors” for major US network carriers on short- and medium-haul routes, often connecting smaller cities to the major carriers’ hubs. Majors subcontract these routes to regional airlines because regionals have a cost advantage in operating small aircraft.⁶ This cost advantage results primarily from the lower salaries and less generous benefits and work rules that regional airline employees receive.⁷ For example, Barry Hirsch (2007) compares the wage rates that United was proposing while in bankruptcy protection in 2004 to prevailing wage rates at 18 regional carriers. He finds that, even after the proposed wage concessions, United’s most senior pilots would be earning 80 percent more than the most senior pilots at the regional airlines. United’s most senior flight attendants and mechanics would be earning 32 percent more and 31 percent more than their counterparts at the regionals, respectively.⁸ During our sample period—which predates the post-2001 bankruptcies and ensuing labor concessions—these comparisons would likely be even larger.

These numbers raise the question of why regional airline employees have been so much less successful at extracting rents. We believe that there are several contributing factors. First, regionals operate in a more competitive labor market. The mobility of employees across regionals is greater than at major carriers, where mobility is limited by union contracts that strictly tie pay to an employee’s seniority with its current employer. Furthermore, regional airlines are able to hire from a pool of individuals willing to work at lower salaries in order to gain the experience

⁴ The incentive problem that arises between majors and independent regionals is similar to that considered by Jackson Nickerson and Brian Silverman (2003) who explore how the externalities that arise in less-than-truckload carriage affect the choice of organizational form.

⁵ For a detailed description of the role of regionals in the US airline industry, see Forbes and Lederman (2007).

⁶ The major network carriers do not operate any small aircraft themselves. Thus, a major’s decision whether to use a regional to serve a particular route is effectively a decision about the type of plane to use for that route.

⁷ Pilots at major airlines recognize the substitutability between regionals’ flights and their own, and negotiate so-called “scope clauses” which limit the extent to which majors can subcontract to regionals. Scope clauses place a limit on the maximum plane size that can be flown by regionals and on the number of such planes that can be flown by regionals. Importantly, scope clauses generally do not distinguish between owned and independent regionals. Thus, the existence of scope clauses should not affect our estimates of the factors affecting a major’s choice of whether to use an owned or independent regional.

⁸ Note that these labor cost differences could not exist if the regional’s employees worked directly for the major. Under the Railway Labor Act, they would have then to be part of the union agreement that exists between the major and its employees, and receive pay, benefits, and work rules consistent with the rest of the major’s employees. While pilots could still be paid based on the type of aircraft they fly, no other craft (i.e., flight attendants, mechanics, ramp workers) has their pay contingent on the types of planes they work on.

necessary to later obtain a job with a larger carrier. Second, union bargaining power is weaker in the regional segment. Not all regional carriers have unionized labor forces. Furthermore, even at regionals that are unionized, the threat of a strike is arguably weaker because the major can, at least in some cases, switch to another regional and thus limit the impact of the strike.

B. Organizational Forms

Almost all regional airlines operate under codeshare agreements with one or more major carriers.⁹ Under these agreements, the regional operates flights on behalf of the major carrier, which markets and tickets these flights under its own flight designator code. In addition to using the major's code, the regional's flights also share the major's brand (for example, Delta's regional Comair operates under the name Delta Connection). Codeshare relationships between major carriers and regionals are governed by one of two types of organizational forms. A regional may be independently owned and contract with one or more major carriers. Or, a regional may be wholly owned by the major with which it partners. This is what we call "vertical integration."¹⁰

Table 1 lists the major-regional partnerships that were in place in 2000 for the large network carriers. These carriers are American Airlines, Continental Airlines, Delta Air Lines, Northwest Airlines, Trans World Airlines, United Airlines and US Airways. Regional carriers that appear in bold were fully owned by their major partner. The table shows that there is substantial heterogeneity both across and within majors in the extent to which regional partners are owned. In 2000, some majors owned all of their regional partners, others owned none, and still others used a mix of owned and independent regional carriers.

C. Contracts and the Need for Adaptation

Relationships between majors and independent regionals are governed by contracts. These contracts generally take one of two forms. Historically, most contracts were structured as revenue-sharing agreements under which the major and regional would share the revenue from passengers whose itineraries involve travel on both airlines.¹¹ More recently, contracts have taken the form of capacity purchase agreements under which the major retains all revenue from flights operated by its regional and pays the regional a fixed fee (usually based on block hours flown) for each departure that the regional operates. Because capacity purchase agreements provide no revenue-based incentives for performance, they often include incentive payments based on operational performance or passenger volumes.

Both types of contracts were in use during our sample period. While we do not have systematic data on which independent regionals were covered by a particular contract form, our empirical analysis does not depend on the type of contract in place. Rather, it depends on two features of these contracts, both of which exist under either contract type. First, contracts between majors and independent regionals are inherently incomplete. While majors and regionals can easily contract on ex ante scheduling decisions (e.g., the set of routes that the regional is to fly, the type of plane that is to be used), majors and regionals cannot contract on ex post changes to the regional's schedule. The set of possible schedule disruptions is so large that it would be prohibitively costly for the major and the regional to specify a contract covering all possible contingencies.

⁹ In 2003, 99 percent of regional airline passengers traveled on flights that were codeshared with a major carrier.

¹⁰ Wholly owned regionals are operated as subsidiaries rather than being fully integrated into the major's operations, so that the regional can legally maintain labor contracts that are distinct from the major's, and thereby preserve the cost advantages that the regional has. See footnote 7 above.

¹¹ The regional would receive 100 percent of the revenue from passengers traveling entirely on the regional.

TABLE 1—MAJORS AND REGIONAL PARTNERS IN 2000

Major	Regional partner
American Airlines	American Eagle Airlines Business Express
Continental Airlines	Continental Express Gulfstream International Airlines
Delta Air Lines	Atlantic Coast Airlines/ACJet Atlantic Southeast Airlines Comair SkyWest Airlines Trans States Airlines
Northwest Airlines	Express Airlines I Mesaba Aviation
Trans World Airlines	Chautauqua Airlines Trans States Airlines
United Airlines	Air Wisconsin Atlantic Coast Airlines Great Lakes Aviation Gulfstream International Airlines SkyWest Airlines
US Airways	Mesa Air Group/Air Midwest Allegheny Airlines Mesa Air Group/CCAir Chautauqua Airlines Colgan Airways Commutair Mesa Air Group/Mesa Airlines Piedmont Airlines PSA Airlines

Note: Regional carriers in bold are fully owned by the major.

Source: Regional Airline Association (www.raa.org).

Moreover, contracting on these contingencies may be undesirable since it could eliminate valuable flexibility.¹² Second, under both types of contracts, independent regionals face financial incentives that are based only on the routes they serve. As a result, when noncontracted schedule adjustments become necessary, majors and independent regionals may disagree on what adjustments should be carried out. In particular, while a major will seek to maximize the overall profitability of its network, an independent regional will seek to maximize its own profits, which are based only on its performance on the narrow set of routes that it serves.¹³ Thus the transaction costs resulting from haggling between majors and independent regionals over ex post changes and adaptations are potentially large.

Of course, one might question whether these transaction costs could not be reduced by contractually allocating control rights to the major. Our understanding from conversations with industry participants is that many contracts between majors and regionals do, in fact, allocate the rights to decide on schedule adjustments to the major. However, having the rights to order

¹² Patrick Bajari and Tadelis (2001) endogenize the level of contractual incompleteness, showing that it may be desirable not to specify complex decisions ex ante in order to preserve flexibility and reduce the costs of ex post negotiations.

¹³ Note that, under both contract types, regionals are the residual claimants on any profits resulting from effective management of labor costs. As a result, they will be particularly averse to schedule changes that impose additional labor costs on them (for example, by forcing delays).

specific schedule changes is not equivalent to having the rights to implement those schedule changes.¹⁴ Schedule changes ordered by the major must still be carried out by the regional; thus, we expect that incentive problems will remain and haggling costs will simply be replaced by the costs of monitoring and enforcing the regional's compliance with the major's requests. The remainder of the paper investigates whether ownership may reduce these costs.

D. *The Costs of Integration*

While ownership of a regional may have the benefit of reducing transaction costs, there are also costs to majors' ownership of their regionals. First, once a major carrier owns a regional and faces transaction costs of divesting it, it may be difficult for the major to commit not to treat the regional's employees symmetrically to their unionized mainline employees.¹⁵ Second, because regionals that are owned by a major do not perform contract flying for competing majors, an owned regional's ability to assemble an efficient route structure may be compromised by its exclusive relationship with a single major.¹⁶ Finally, as in all situations in which a firm faces an integration decision, nonintegration may have the effect of sharpening performance incentives. However, this is unlikely to be a major factor in our setting because the fixed-fee contracts that are increasingly being used leave an independent regional with no residual incentives to increase demand (although they preserve the incentive to cut costs). While we do not observe these cost differences directly, we assume that they do not vary with city pair characteristics.¹⁷

II. Empirical Approach

To investigate whether major airlines choose ownership when they would face more frequent or more costly negotiations with independent regionals over ex post adaptations, we exploit the fact that, in our setting, both the frequency of adaptation decisions and the costliness of having such adaptation decisions resolved in a less-than-optimal way vary with observable market characteristics. First, the frequency of adaptation decisions will be affected by the weather patterns at the endpoint airports of each city pair. Adverse weather conditions lead to unanticipated schedule adjustments by increasing the amount of time that is needed in between consecutive takeoffs or landings, thus forcing airlines to delay or cancel flights. Indeed, weather is the leading cause of flight delays that is not under the airline's control.

Second, both the frequency of adaptation decisions and the costliness of having adaptation decisions resolved in a less-than-optimal way are affected by the extent to which a city pair is integrated into an airline's network. The more integrated a city pair, the more likely it is that flights on that city pair will be affected by unanticipated disruptions that occur elsewhere in the network. Furthermore, the more integrated a city pair, the more likely it is that flights on that city pair will impose externalities on other flights in the airline's network. This is both because these

¹⁴ This is in contrast to the standard Grossman-Hart depiction of control rights, which confer the right to determine precisely how an asset is used.

¹⁵ For example, after Delta purchased the previously independent regional Comair, Comair's pilots went on strike demanding compensation comparable with that of Delta's pilots. Similarly, in response to an announcement in 2007 that American Airlines was considering spinning off American Eagle, American Eagle union representatives voiced concern that the spin-off would result in lower compensation for Eagle employees.

¹⁶ Technically, it would be possible for an owned regional to fly under contract for a competing major. However, one would expect that the potential for a hold-up problem would deter competing majors from entering such a relationship. Empirically, we never observe owned regionals flying under contract for competing carriers.

¹⁷ Insofar as they result in higher labor costs at owned regionals, this should be constant across city pairs after controlling for distance because pilots and flight attendants are paid on block hours.

flights will have to compete with the airline's other flights for access to scarce airline and airport facilities and because these flights are more likely to carry connecting passengers.

Our empirical analysis tests whether a city pair's weather conditions and its degree of integration into the major's network affect the major's choice to vertically integrate with its regional partner on the city pair. However, we recognize that the major also has the option to serve the city pair with its own planes, and therefore we estimate a nested logit model, and we assume that the major first chooses whether to serve a city pair itself or with a regional and then, conditional on using a regional carrier, chooses whether to use an owned or independent one.

One advantageous feature of our data is that we have four airlines that use both owned and independent regionals during our sample quarter. Therefore, in some specifications, we include only the four major carriers that used both types of regional partners in 2000. In these specifications, we include in both the top and bottom nests a dummy variable for each major carrier to control for that carrier's unobserved propensity to use regionals and its unobserved propensity to use regionals of a given type.

The key identifying assumption of our empirical approach is that none of our explanatory variables is correlated with other unobserved factors that affect the relative returns to serving a city pair in one of the three possible ways. For example, we assume that, after controlling for distance, any unobserved operating cost differences between owned and independent regionals are not correlated with the network and weather variables. Finally, we consider the network structure and, in particular, hub locations to be predetermined.

III. Data and Measurement

A. Data and Sample¹⁸

Data from the Official Airline Guide (OAG) provide the complete flight schedules of all domestic airlines. This data allow us to identify every flight that a major subcontracts to a regional, as well as the identity of that regional. We combine this with data from the Regional Airline Association (RAA) which identify the ownership type of each regional. Finally, weather data for each airport location are taken from the National Oceanic and Atmospheric Administration (NOAA). These data include monthly observations over 25 years (1971–1995) on precipitation and average daily minimum temperature (average taken over all days of the month), as well as a separate source of data on average annual snowfall (average taken over 30 years, 1971–2000).

Our sample period is the second quarter of 2000. We consider flights between the top 300 US airports.¹⁹ Our level of observation is the airline-city pair. We use the term "city pair" to refer to direct nonstop service between two endpoint airports, in either direction. We do not distinguish the direction of service because airlines do not make independent decisions on which organizational form to use in each direction. In some cases, a major will offer several flights per day on a city pair and will operate some set of those flights itself and others with a regional partner. For these city pairs, we investigate the type of regional partner that is used for the flights operated by the regional. Our sample is limited to city pairs with a distance of under 1,500 miles. These are city pairs that could be served by a regional with one of its planes.

¹⁸ Data and programs for the estimation are available in the Web Appendix (<http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.5.1831>).

¹⁹ Airport rankings are based on total passenger enplanements and are available at <http://www.faa.gov/arp/planning/stats/2001/prim01.xls>. We exclude routes that have an endpoint in Alaska, Hawaii, Puerto Rico, Guam, or the US Virgin Islands. The smallest airport in our sample is Grand Canyon National Park Airport, with two flights per day.

TABLE 2—VARIABLE NAMES AND DEFINITIONS

Variable name	Definition	Source
Regional	Equals 1 if flight is operated by a regional of either type	OAG
Owned regional	Equals 1 if flight is operated by an owned regional	OAG and RAA
Hub	Equals 1 if either endpoint is carrier's hub	Authors' construction
Number of flights at larger endpoint	Carrier's number of departing domestic flights per day on other city pairs from endpoint at which it is larger, in hundreds	OAG
Number of flights at smaller endpoint	Carrier's number of departing domestic flights per day on other city pairs from endpoint at which it is smaller, in hundreds	OAG
Precipitation	Average annual precipitation, in inches. We use the maximum of these averages across the two endpoint airports of a city pair (based on 1971–1995 data). For freezing or frozen precipitation, we measure the water equivalent.	NOAA
Snowfall	Average annual snowfall, in inches. We use the maximum of these averages taken across the two endpoint airports of a city pair (based on 1971–2000 data).	NOAA
Number of freezing months per year	Average number of months per year on which average minimum temperature is below 0 degrees Celsius. We use the maximum of these averages taken across the two endpoint airports of a city pair (based on 1971–1995 data).	NOAA
Distance	Distance of the city pair, in hundreds of miles	Authors' construction
Slot	Equals 1 if either endpoint is slot controlled	Authors' construction
Population at larger endpoint	Population of the Metropolitan Statistical Area in the year 1999 at the endpoint with the larger population, in thousands	US Census Bureau
Population at smaller endpoint	Population of the Metropolitan Statistical Area in the year 1999 at the endpoint with the smaller population, in thousands	US Census Bureau

We consider the organizational form decisions of the seven large network carriers (shown in Table 1). We exclude codeshare relationships between majors and regionals in which the major sells tickets for seats on the regional under its own code, but in which the regional does not operate under the name and brand of the major. We restrict our analysis to flights that operate on Mondays, thus considering only the weekday schedules of the airlines.²⁰ After imposing these restrictions and eliminating observations with missing values for any of our explanatory variables, we arrive at a sample of 1,745 airline-city pairs: 56 percent of these are served by the major itself, 26 percent are served by owned regionals, and the remaining 19 percent are served by independent regionals.

B. Variables

Variable names and definitions appear in Table 2. Summary statistics appear in Table 3.

Variables That Influence the Type of Regional Used.—Our main variables of interest in the bottom nest are those that measure the frequency and costliness of adaptation decisions. We construct two measures of the extent to which a city pair is integrated into the major's network. The dummy variable *Hub* equals one if either endpoint of the city pair is the major's hub. On airline-city pairs that involve a hub, 60 percent of all passengers connect to or have connected

²⁰ Of the flights that operate on Mondays, 99.46 percent also operate Tuesday through Friday.

TABLE 3—SUMMARY STATISTICS

Variable name	N	Mean	SD	Min	Max
Regional	1,742	0.44	0.50	0	1
Owned regional, given Regional equals 1	774	0.58	0.49	0	1
Hub	1,742	0.76	0.43	0	1
Number of flights at larger endpoint (in hundreds)	1,742	3.55	2.12	0.02	8.23
Number of flights at smaller endpoint (in hundreds)	1,742	0.19	0.41	0	4.91
Precipitation	1,742	37.27	8.73	6.45	55.85
Snowfall	1,742	33.44	25.08	0	118.1
Number of freezing months per year	1,742	3.26	1.58	0	6.27
Distance (in hundreds of miles)	1,742	5.57	3.63	0.20	14.99
Slot	1,742	0.20	0.40	0	1
Population at larger endpoint (in thousands)	1,742	6,744.24	5,811.34	89.34	20,196.65
Population at smaller endpoint (in thousands)	1,742	1,639.23	1,592.54	63.16	8,885.92

from another flight, compared to 26 percent on airline-city pairs that do not involve a hub.²¹ Although, regionals are largely used to provide feeder traffic to majors' hubs, 29 percent of city pairs served by regionals in our sample do not have the major's hub on either endpoint. It is these routes that provide the variation used to identify the *Hub* variable.

Our second measure of a city pair's network integration is the total number of flights operated by the major or its regionals from the endpoint airports of a city pair on a given day.²² These variables proxy for the number of flight connections that can be made. The measures allow us to investigate the relationship between an airline's network at an airport and organizational form using a subsample that includes only nonhub airports. It also provides a continuous measure of network size. We call these variables *Number of flights at larger endpoint* and *Number of flights at smaller endpoint*.

We construct three variables that measure expected weather conditions. First, for each endpoint airport, we calculate the average annual precipitation.²³ We then take the maximum of this variable across the two endpoint airports of a city pair and call this *Precipitation*. We follow the same procedure to construct *Snowfall*. Finally, to ensure that our snowfall variable is capturing the impact of snow and not cold weather, we calculate the average number of months per year in which the average daily minimum temperature at an airport is below 32 degrees Fahrenheit. We then take the maximum of this variable across the two endpoint airports of a city pair and call this *Number of freezing months per year*.

We also include several control variables in the bottom nest. *Slot* is a dummy variable that is equal to one if either endpoint of the city pair is one of four slot-controlled airports.²⁴ Slot controlled airports are highly congested and prone to flight delays, and airlines may need to make more rescheduling decisions at these airports. Our second control variable is the distance of the city pair (*Distance*). On longer city pairs, the labor cost advantage that independent regionals have over owned regionals may be greater because pilots and flight attendants are compensated based on block hours.

²¹ These estimates are based on passenger numbers from the Department of Transportation DB1A database for airline-city pairs in our sample.

²² This measure does not include flights on the city pair being considered.

²³ This average is taken over the 25 years of monthly weather data that we have. When precipitation is frozen (i.e., snow, hail, or freezing rain), this variable measures the water equivalent of the precipitation, which is different from the depth of snowfall. The density of new snow is typically between 5 percent and 12 percent of water.

²⁴ These airports are Chicago O'Hare, John F. Kennedy and LaGuardia Airports in New York, and Reagan National in Washington, DC.

Variables That Influence Whether to Use a Regional (of Either Type).—A major's decision whether to serve a city pair itself or subcontract that city pair to a regional is effectively a decision about plane size. Several factors will determine the suitability of a route for service by a small aircraft. The first and probably most important is expected demand. To proxy for the expected demand on a city pair, we construct *Population at larger endpoint* and *Population at smaller endpoint* which measure the population of the Metropolitan Statistical Area at the endpoint with the larger and smaller population, respectively. We expect that city pairs that involve one endpoint with a very small population are more likely to be served by a regional than by the major itself.²⁵ We also include *Hub* in the top nest. Controlling for endpoint population, hub airports may have higher demand than nonhub airports.

A second factor that will affect the optimal aircraft size on a given city pair is the *Distance* of the city pair. Since regional aircraft are generally considered to be less comfortable and noisier than larger jets, they are typically used on short routes. Regional aircraft also tend to be more efficient on shorter routes. Thus, we expect that longer routes are more likely to be served by the major itself than by a regional.

An additional factor that might affect a major's choice between large and small planes is whether one of the endpoint airports is slot controlled. One reason to use small planes on a route is to offer high-frequency service. At airports that are slot controlled, however, the number of takeoffs and landings is restricted. It is generally more profitable to use these takeoff and landing slots for larger planes. We therefore include *Slot* in the upper-level nest and expect it to have a negative effect on the likelihood of a regional being used. We do not include the weather variables in the top nest since we have no reason to expect that weather conditions affect the relative returns to using larger versus smaller aircraft.

IV. Results

A. Estimation Results

Tables 4 and 5 present the results of our empirical analysis. The tables are divided into two panels. The upper panel presents the coefficients on variables that influence a major's decision between using an owned and independent regional (the bottom nest), while the lower panel presents the coefficients on the variables that influence a major's decision whether to serve a city pair itself or with a regional of either type (the top nest). Both tables present coefficient estimates.

Column 4-1 presents our baseline specification. In the bottom nest, we include the *Hub* variable, the three weather variables (*Precipitation*, *Snowfall*, and *Number of freezing months per year*), as well as the control variables *Distance* and *Slot*. The results indicate that city pairs with the major's hub at either endpoint are significantly more likely to be served by a regional that the major owns. Holding all other variables at their means, the coefficient estimate implies that having the major's hub on either endpoint increases the conditional probability of using an owned regional by 57.4 percentage points, compared to having no hub at either endpoint.²⁶

Further, the estimates in 4-1 show that city pairs with higher average precipitation are more likely to be served by an owned regional. This is also true for city pairs that experience more snowfall. These initial findings suggest that owned regional carriers are more likely to be used

²⁵ Note that we do not include the population variables in the bottom nest since these variables should have no effect on either the costs or benefits of using a wholly owned regional.

²⁶ When we calculate the marginal effect on the conditional probability of using an owned regional, we hold all other variables at their means in the subsample of data that includes only city pairs on which a major uses a regional of either type. This is true for all marginal effects that we calculate in this section for variables that appear in the bottom nest. For dummy variables, we report the effect of changing the value of the dummy from zero to one.

TABLE 4—NESTED LOGIT

	(4-1)	(4-2)	(4-3)	(4-4)
<i>Panel A. Bottom nest</i>				
	Dependent variable equals 1 if owned regional is used			
Hub	1.429 (0.205)***	1.012 (0.280)***	3.527 (0.381)***	3.571 (0.440)***
Precipitation	0.022 (0.005)***	0.004 (0.005)	0.016 (0.006)**	0.002 (0.006)
Snowfall	0.020 (0.005)***	0.021 (0.005)***	0.024 (0.006)***	0.021 (0.006)***
Number of freezing months per year	-0.711 (0.087)***	-0.737 (0.091)***	-0.756 (0.123)***	-0.735 (0.124)***
Distance	0.030 (0.040)	0.035 (0.040)	-0.122 (0.063)*	-0.090 (0.062)
Slot	0.640 (0.222)***	0.600 (0.225)**	1.398 (0.370)***	1.206 (0.384)***
Number of flights at larger endpoint		0.230 (0.063)***		0.025 (0.072)
Number of flights at smaller endpoint		2.229 (0.457)***		2.387 (0.548)***
<i>Panel B. Top nest</i>				
	Dependent variable equals 1 if any regional is used			
Hub	0.063 (0.162)	0.279 (0.169)	1.055 (0.385)***	1.526 (0.395)***
ln (population at larger endpoint)	0.869 (0.064)***	0.850 (0.065)***	0.810 (0.081)***	0.747 (0.082)***
ln (population at smaller endpoint)	-0.724 (0.071)***	-0.672 (0.073)***	-0.615 (0.089)***	-0.513 (0.092)***
Distance	-0.434 (0.025)***	-0.441 (0.028)***	-0.513 (0.042)***	-0.529 (0.045)***
Slot	-0.860 (0.174)***	-0.770 (0.191)***	-0.395 (0.300)	-0.082 (0.326)
Inclusive Values				
Regional	-0.250 (0.130)*	-0.508 (0.111)***	-0.603 (0.163)***	-0.738 (0.164)***
Carrier dummies	No	No	Yes	Yes
Observations	1,742	1,742	1,159	1,159

Notes: Standard errors in parentheses. Sample includes city pairs between top 300 US airports operated by major network carriers (or their regional partners) in the second quarter of 2000. Columns 3 and 4 include only Continental, Delta, Northwest, or US Airways (or their regional partners). *Precipitation* and *Snowfall* are measured in inches. *Distance* is measured in hundreds of miles. *Number of flights at larger (smaller) endpoint* is measured in hundreds. *Population at larger (smaller) endpoint* is measured in thousands.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

on city pairs on which adaptation decisions are more likely to be needed. At the sample means, an additional 0.1 inch of *Precipitation* would increase the conditional probability of using an owned regional by 0.11 percentage points. Assuming an average water equivalent of 8 percent, the estimate on the *Snowfall* variable implies that an increase in the depth of snowfall of 1.25 inches, or 0.1 inch of water equivalent, would increase the conditional probability of using an owned regional by 1.27 percentage points.

The coefficient estimate on *Number of freezing months per year* indicates that cold temperatures decrease the likelihood that a major uses an owned regional. If weather indeed influences vertical integration decisions by affecting the likelihood of adaptation decisions, then the

TABLE 5—NESTED LOGIT (*All Carriers*) HUB ROUTES VERSUS NON-HUB ROUTES

	(5-1)	(5-2)	(5-3)	(5-4)
	Hub routes	Hub routes	Non-hub routes	Non-hub routes
<i>Panel A. Bottom nest</i>				
Dependent variable equals 1 if owned regional is used				
Precipitation	0.087 (0.009)***	0.058 (0.010)***	-0.018 (0.003)***	-0.010 (0.008)
Snowfall	0.021 (0.007)***	0.027 (0.007)***	0.001 (0.001)	0.022 (0.008)***
Number of freezing months per year	-0.891 (0.120)***	-0.980 (0.124)***	-0.026 (0.024)	-0.488 (0.159)***
Distance	-0.011 (0.050)	-0.004 (0.049)	0.181 (0.039)***	-0.117 (0.084)
Slot	0.363 (0.318)	0.314 (0.327)	0.150 (0.231)	0.254 (0.358)
Number of flights at larger endpoint		0.186 (0.072)**		1.371 (0.315)***
Number of flights at smaller endpoint		5.472 (1.202)***		-1.076 (0.677)
<i>Panel B. Top nest</i>				
Dependent variable equals 1 if any regional is used				
ln (population at larger endpoint)	0.828 (0.070)***	0.817 (0.071)***	1.337 (0.206)***	0.780 (0.172)***
ln (population at smaller endpoint)	-0.699 (0.081)***	-0.619 (0.084)***	-0.688 (0.173)***	-0.722 (0.167)***
Distance	-0.422 (0.028)***	-0.435 (0.030)***	-0.526 (0.348)	-0.478 (0.066)***
Slot	-1.447 (0.222)***	-1.514 (0.232)***	0.042 (1.447)	0.013 (0.344)
Inclusive values				
Regional	-0.064 (0.073)	-0.264 (0.060)***	-11.836 (2.506)***	0.955 (0.670)
Observations	1,317	1,317	425	425

Notes: Standard errors in parentheses. Sample includes city pairs between top 300 US airports operated by the major network carriers (or their regional partners) in the second quarter of 2000. Columns 1 and 2 include city pairs with airline's hub on either endpoint. Columns 3 and 4 include city pairs without airline's hub on either endpoint. *Precipitation* and *Snowfall* are measured in inches. *Distance* is measured in hundreds of miles. *Number of flights at larger (smaller) endpoint* is measured in hundreds. *Population at larger (smaller) endpoint* is measured in thousands.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

estimate on *Number of freezing months per year* suggests that in cold weather such adaptation decisions are less likely. To confirm this, we investigate the relationship between weather and *actual* flight delays. We describe the data and sample used for this analysis in more detail in the Appendix (further materials are available in the Web Appendix). The results, presented in Table A1, indicate that flight delays *increase* with the amount of daily precipitation and, especially, with precipitation on days with below freezing temperatures.²⁷ Consistent with our finding in column 4-1, below freezing temperatures—on their own—have a *negative* impact on flight delays.

²⁷ Our snowfall data are available only as a 30-year average and not at a disaggregated level; therefore, in the regressions in the Appendix, we use precipitation on days with below freezing temperature instead.

Note that this relationship is true even when we include airline-city pair fixed effects. Thus, the cold weather variable is not capturing some unobservable characteristic that varies across airports in a way that is correlated with weather patterns. Rather, we believe that the cold weather measure is capturing the fact that very cold days tend to be very clear days and therefore quite well suited for flying.²⁸ Importantly, the results in the Appendix indicate that the relationship we estimate between weather and organizational form is consistent with the relationship we estimate between weather and actual flight delays.

With respect to our additional control variables, we find that the coefficient on *Distance* is insignificant. Owned and independent regionals in our sample serve routes of very similar distance, making any possible effect of *Distance* on the decision to vertically integrate difficult to identify. *Slot* has a positive effect on the likelihood of using an owned regional. Since schedule disruptions are more likely at slot constrained airports, this provides additional evidence that majors appear to use owned regionals when adaptation decisions need to be made more frequently.

We now briefly discuss the estimates for the variables included in the top nest. We find that increasing the population of the larger endpoint airport of a city pair increases the likelihood of using a regional, while increasing the population of the smaller endpoint decreases the likelihood of using a regional. These estimates suggest that city pairs connecting a large city with a small city are the ones that are most likely to be served by a regional. *Distance* has a negative and significant coefficient in the top nest, indicating the longer city pairs are less likely to be served by a regional. *Slot* also has a negative and significant coefficient, indicating that slot controlled airports (at which takeoff and landing slots are harder to secure) are less likely to be served by small aircraft. Finally, the estimate on the hub variable is insignificant.

In 4–2, we add our continuous network variables, *Number of flights at larger endpoint* and *Number of flights at smaller endpoint* to the bottom nest. The coefficients on these variables in 4–2 show that, even when controlling for *Hub*, the number of flights that the major operates at both the larger and the smaller endpoint has a positive effect on the likelihood of using an owned regional. Adding an additional flight at the smaller endpoint has a much larger effect than adding one flight at the larger endpoint, which is the endpoint that already has more flights to which connections can be made. The relative magnitudes of the marginal effects imply that one additional flight at the smaller endpoint has the same effect on the conditional probability of using an owned regional as nine additional flights at the larger endpoint.

The coefficient on *Precipitation* turns insignificant in this specification, but the coefficient on *Snowfall* remains significant and maintains the same magnitude as in 4–1. Snowfall is much more likely than precipitation to cause schedule disruptions, and this specification, together with the relative magnitudes of the marginal effects computed for specification 4–1 above, suggests that the decision to use an owned regional on a given city pair is much more strongly influenced by the annual amount of snowfall at the endpoints of the city pair than by the annual amount of rain.

Next, we restrict the sample to the major airlines that use both types of regionals and include a dummy variable for each of these major carriers. We reestimate 4–1 and 4–2 on this reduced sample and report the results in the third and fourth columns of the table. We find in these specifications that the estimated effects of the network and weather variables have the same signs as in the larger sample and remain statistically significant. The estimates of the top nest are quite similar in the restricted sample. The only differences from the previous results are that, in the restricted sample, *Hub* is estimated to have a positive effect on the likelihood of using a regional,

²⁸ Note that Michael Mazzeo (2003) also finds a negative relationship between below freezing temperatures and flight delays.

implying that these majors use regionals extensively at their hubs, while *Slot* has no statistically significant impact.

In Table 5, we split the sample into city pairs that have the carrier's hub at one of the endpoints—we call these *hub routes*—and city pairs that do not have the carrier's hub at either endpoint—which we call *nonhub routes*. We reestimate specifications 4-1 and 4-2 on these two subsamples, with the hub variable, of course, omitted. This allows us to investigate whether our other explanatory variables affect hub and nonhub city pairs in similar ways. However, because over 75 percent of the routes in our sample are hub routes, the sample size for the nonhub regressions is small. Looking first at hub routes, the results are quite similar to those in Table 4. Interestingly, we find in 5-2 that—even considering hub routes only—variation in the extent to which a city pair is integrated into the major's network, as measured by the number of flights at the endpoint, still affects the likelihood of using an owned regional.

Columns 5-3 and 5-4 present the results for nonhub routes. In 5-3 the coefficient on *Precipitation* turns negative and *Snowfall* and cold weather are no longer significant. Once we add our continuous measures of network size in 5-4, the coefficient on *Precipitation* is no longer significant. *Snowfall* again increases the likelihood of using an owned regional, and cold weather still has a strongly negative effect. *Distance*, *Slot*, and the number of flights at the smaller endpoint have no effect, but the number of flights at the larger endpoint has a positive and significant effect on the likelihood of using an owned regional. This suggests that, on nonhub routes, what matters is whether the route involves one endpoint airport at which the airline has a reasonably large number of flights.²⁹ In all four specifications in Table 5, the top nest results are very similar to those in the previous table.

B. Robustness Checks

We performed several robustness checks on our results. First, we used an alternative way of measuring the weather across the two endpoint airports of a city pair. We calculated the average of each weather measure across the two endpoints, instead of the maximum. The results are robust to this variation on the weather measures.³⁰ Second, we used nonlinear functional forms for distance, as well as a continuous measure of airport congestion constructed as the number of daily domestic flights divided by the number of parallel runways at the airport. Our results are robust to these changes.

Next, we address the issue that *Snowfall* might proxy for the unionization of regional carriers. This is because the Northeast, the region with the largest amount of snowfall, is more unionized than the rest of the United States. In highly unionized states, the cost advantage of independent regionals (with strong unions) may be smaller than in other states where unions are weaker. This may make the use of independent regionals less attractive. To investigate this possibility, we determine for each regional in our sample the state in which it has its headquarters. Henry Farber (1984) shows that states with “right-to-work” laws have preferences against union representation. There are 22 right-to-work states, most of which are in the South and in the West. We include a dummy variable in our regressions that controls for the regional's headquarters being located in

²⁹ There are several nonhub airports at which airlines still operate a reasonably large number of flights—for example, US Airways in Boston and Delta in Orlando. The results in 5-4 suggest that what matters for the vertical integration decision—for nonhub routes—is whether the route involves one of these types of airports. The size of the carrier at the smaller airport appears to be irrelevant.

³⁰ Note that the average length of routes served by regionals in our sample is 337 miles. We would expect not to have very large differences in weather patterns across the two endpoints. Thus, it is not surprising that using the average weather and the maximum weather produces consistent results.

a more unionized (non-right-to-work) state.³¹ Our results are robust to including this control in our regressions.

C. Airport Level Analysis

The results in the previous section relate a major's decision to use an owned regional on a particular city pair to characteristics of that city pair. However, given that majors often serve a large number of city pairs out of a given airport, one might question whether majors actually decide which type of regional to use at a particular airport and then use that same regional for all regional city pairs out of that airport. In particular, one could imagine that majors might determine which regional to use at each of their hubs (and perhaps at other airports at which they have reasonably sized network operations) and then use that regional for all routes that connect to that airport. Note that this would be likely, for example, if there were economies of scope in having the same regional serve multiple city pairs out of the same airport, or if there were contracting complementarities across those city pairs.

We believe that, in practice, majors do decide which type of regional to use on a city pair-by-city pair basis. Majors often use multiple regionals, and even regionals of both types, at the same airport.³² In Tables 6A and 6B, we investigate majors' patterns of regional usage at the airport level. We create a dataset that includes all airline-airport combinations at which one of our seven majors serves at least two regional routes (so that the decision to use more than one regional carrier is relevant). In Table 6A, we document the frequency with which majors use multiple regionals at the same airport. We find that, in our sample, the *same* major uses multiple regionals at an airport for 124 of our 413 carrier-airport observations, or 30 percent of the time. To ensure that this figure is not simply capturing spoke airports at which the major uses different regionals to connect to different hubs, we also look specifically at hub airports and at what we call "large" airports (which we define as airports at which an airline operates more than 35 flights per day). Looking at the set of 23 airline-airport observations that are hubs, we find that majors use multiple regionals at seven of these (or 30 percent of the time). They use multiple regionals at 19 of the "large" airports in our sample (or 45 percent of the time).

In Table 6B, we consider just the four major carriers that use both owned and independent regionals. We examine the extent to which these four carriers use both types of regionals at a given airport, again distinguishing between all airports and hubs (as well as "large" airports). We find 73 instances of one of these four majors using both types of regionals at the same airport (26 percent of the time). Furthermore, these majors use both types of regionals at 30 percent of hubs and at 50 percent of "large" airports. Thus, the evidence in Tables 6A and 6B—combined with anecdotal evidence from industry participants—suggests quite strongly that majors do not (and certainly need not) choose a single regional or a single type of regional to serve all routes out of a given airport.

D. Alternate Explanations

Before concluding, we briefly consider whether a major might want to own a regional carrier in order to limit its rivals' access to that regional (i.e., if foreclosure might be a motive for integration). We think that this is unlikely because barriers to entry in the regional airline market are

³¹ Even though right-to-work laws do not apply to employees in the airline industry, who are governed by the Railway Labor Act, the existence of such laws reflects the tendency toward unionization in the state.

³² We also asked industry participants if there are reasons to use the same regional for all city pairs served out of an airport; we were consistently told that there are not.

TABLE 6A—PATTERNS OF REGIONAL USAGE, AIRPORT ANALYSIS—ALL MAJORS

	All airports	Hubs	“Large” airports	“Small” airports
Observations	413	23	42	371
Number at which major uses more than one regional	124 (30.02 percent)	7 (30.43 percent)	19 (45.24 percent)	105 (28.3 percent)
Number at which major uses more than two regionals	35 (8.47 percent)	4 (17.39 percent)	9 (21.43 percent)	26 (7.01 percent)

Notes: Level of observation is the airline-airport. Sample includes airports at which a major serves more than one regional route.

TABLE 6B—PATTERNS OF REGIONAL USAGE, AIRPORT ANALYSIS—MAJORS WITH BOTH TYPES OF REGIONALS

	All airports	Hubs	“Large” airports	“Small” airports
Observations	276	13	26	250
Number at which major uses both types of regional	73 (26.45 percent)	4 (30.77 percent)	13 (50 percent)	60 (24 percent)

Notes: Level of observation is the airline-airport. Sample includes airports at which Continental, Delta, Northwest, or US Airways serves more than one regional route.

quite low. As Table 1 indicates, there is a large number of independent regionals, many of which operate for multiple majors. Their assets are highly mobile and many of these regionals already operate in several different parts of the country. This suggests that regionals can easily initiate contract service for a new major, even in geographic areas they have not previously served. As a result, while ownership of a particular regional may prevent rivals from contracting with *that* regional, it is unlikely to prevent rivals from contracting with *any* regional.

V. Conclusion

We have examined if the likelihood and costliness of ex post adaptation decisions affects firms' incentives to vertically integrate. Our results indicate the existence of a strong empirical relationship. We have found that majors are more likely to use owned regionals on city pairs with adverse weather and on city pairs that are more highly integrated into their network. Importantly, we find qualitatively similar results when we restrict our sample to the set of majors that use both owned and independent regionals and include firm dummies in our regressions. We interpret our results as suggesting that ownership allows firms to reduce the transaction costs that arise when negotiations over ex post adaptation decisions become necessary.

While our empirical analysis is carried out for a specific industry, we believe that its implications are much broader. Our results show that externalities across transactions—in this case resulting from the organization of transactions in a network—can be an important factor influencing firms' boundary decisions. In other settings as well, firms seeking to subcontract a portion of their network may find it difficult to use contracts to provide subcontractors with incentives to internalize the impact of their actions on the remainder of their network. In addition, the industry we study is characterized by substantial bargaining power on the part of labor, as well

as a fairly regulated labor market. The presence of unions, as well as the anticipated practices of the National Mediation Board, give rise to costs that exist only when transactions are organized internally. Thus, our analysis suggests that market power in input markets may be an additional factor shaping firms' boundaries.

APPENDIX: THE LINK BETWEEN WEATHER AND FLIGHT DELAYS

To assess if our measures of the expected, or typical, weather at the endpoint airports indeed capture the frequency of adaptation decisions that need to be made on a given city pair, we investigate the relationship between actual weather and actual flight delays. While we do not have information on flight delays of regional carriers, major airlines are required to report flight delays on a flight-by-flight basis to the Bureau of Transportation Statistics. We use these data for the years 1998–2000 to establish the relationship between flight delays and the weather measures we use in the estimation above. Since our data cover only flights by large planes, which are always operated by major airlines, we do not need to account for selection bias coming from the choice of ownership, as we would have to do if we studied this relationship for regional airlines.

We have daily data on the precipitation and the average minimum temperature for each airport in our sample for the years 1998–2000. For each day, we compute average arrival delays for all airline-city pairs between the airports in our sample that are served by major airlines. Arrival delays are defined as the difference between actual and scheduled arrival time, in minutes. We count early arrivals as negative delays. Similar results are obtained when early arrivals are defined as zero delays. In Table A1, we regress the average flight delay on a given day for a given airline on a given city pair on the weather at the endpoint airports on that day, using ordinary least squares (OLS). We measure the weather at the endpoint airports by the maximum precipitation across the two endpoints on that day, measured in inches (*Precipitation*), a dummy for whether either of the endpoints had below freezing temperatures on that day (*Below Freezing*), and the maximum across the two endpoints of the interaction of *Precipitation* and a dummy for the temperature being below freezing at that endpoint. We use this interaction term to proxy for snowfall because we do not have daily snowfall data. We also include a dummy for either of the two endpoints being the carrier's hub.³³

We present four specifications. The first specification in A1-1 regresses flight delays only on the weather variables and the hub dummy. Next, we include carrier fixed effects in A1-2, then city pair fixed effects in A1-3, and finally carrier-city pair fixed effects in A1-4³⁴. The results are strikingly similar across these four specifications. Even after controlling for carrier-city pair fixed effects, i.e., after eliminating all variation that is common for a carrier on a given city pair over the three years of our sample, we find that flight delays are significantly longer on days with precipitation and, in particular, when precipitation occurs on a day with below freezing temperatures. Controlling for precipitation, however, days with below freezing temperatures have significantly shorter flight delays, as compared to days with above freezing temperatures *on the same city pair and for the same carrier*.

These results imply that schedule disruptions are more common on city pairs with more precipitation, and in particular with more snowfall, but are less common on city pairs that have more days with below freezing temperatures. We believe that the cold weather measure is capturing the fact that cold days without precipitation tend to be very clear days and therefore quite well suited for flying.

³³ Chris Mayer and Todd Sinai (2003) find that airlines tend to have longer flight delays at their hubs.

³⁴ This final specification does not include the hub dummy because the effect is not identified when carrier-city pair fixed effects are included.

TABLE A1—THE EFFECT OF WEATHER ON FLIGHT DELAYS

	Dependent variable = average arrival delay on the day			
	(A1-1)	(A1-2)	(A1-3)	(A1-4)
Precipitation	0.1097 (0.0008)***	0.1106 (0.0008)***	0.1107 (0.0008)***	0.1107 (0.0008)***
Below freezing	-0.5537 (0.0433)***	-0.6633 (0.0435)***	-0.7650 (0.0440)***	-0.7715 (0.0440)***
Below freezing × precipitation	0.2649 (0.0049)***	0.2650 (0.0049)***	0.2631 (0.0049)***	0.2631 (0.0049)***
Hub	0.8398 (0.0516)***	0.4634 (0.0530)***	2.0204 (0.1557)***	— —
Fixed effects	None	Carrier	City pair	Carrier-city pair
Observations	1,520,703	1,520,703	1,520,703	1,520,703

Notes: All models estimated with OLS. Heteroskedasticity-robust standard errors in parentheses. Sample includes daily observations for city pairs between top 300 US airports operated by major network carriers, but not regionals, in 1998–2000. *Dependent variable* is average arrival delay for the carrier on the route on that day. *Precipitation* is total precipitation for the day and is measured in hundredths of inches. *Below Freezing* is a dummy that equals one if the minimum temperature on that day was below 0 degrees Celsius.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

REFERENCES

- Anderson, Erin, and David C. Schmittlein. 1984. "Integration of the Sales Force: An Empirical Examination." *RAND Journal of Economics*, 15(3): 385–95.
- Bajari, Patrick, and Steven Tadelis. 2001. "Incentives versus Transaction Costs: A Theory of Procurement Contracts." *RAND Journal of Economics*, 32(3): 387–407.
- Farber, Henry S. 1984. "Right-to-Work Laws and the Extent of Unionization." *Journal of Labor Economics*, 2(3): 319–52.
- Forbes, Silke J., and Mara Lederman. 2007. "The Role of Regional Airlines in the US Airline Industry." In *Advances in Airline Economics*. Vol. II, ed. Darin Lee, 193–208. Amsterdam: Elsevier.
- Gibbons, Robert. 2005. "Four Formal(izable) Theories of the Firm?" *Journal of Economic Behavior and Organization*, 58(2): 200–45.
- Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy*, 94(4): 691–719.
- Hart, Oliver, and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy*, 98(6): 1119–58.
- Hirsch, Barry. 2007. "Wage Determination in the US Airline Industry: Union Power under Product Market Constraints." In *Advances in Airline Economics*. Vol. II, ed. Darin Lee, 27–60. Amsterdam: Elsevier.
- Holmström, Bengt. 1999. "The Firm as a Subeconomy." *Journal of Law, Economics, and Organization*, 15(1): 74–102.
- Holmström, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization*, 7: 24–52.
- Hubbard, Thomas N. 2001. "Contractual Form and Market Thickness in Trucking." *RAND Journal of Economics*, 32(2): 369–86.
- Joskow, Paul L. 1985. "Vertical Integration and Long-Term Contracts: The Case of Coal-burning Electric Generating Plants." *Journal of Law, Economics, and Organization*, 1(1): 33–80.
- Lafontaine, Francine, and Margaret Slade. 2007. "Vertical Integration and Firm Boundaries: The Evidence." *Journal of Economic Literature*, 45(3): 629–85.
- Masten, Scott E. 1984. "The Organization of Production: Evidence from the Aerospace Industry." *Journal of Law and Economics*, 27(2): 403–17.
- Mayer, Christopher, and Todd Sinai. 2003. "Network Effects, Congestion Externalities, and Air Traffic Delays: Or Why Not All Delays Are Evil." *American Economic Review*, 93(4): 1194–1215.

- Mazzeo, Michael J.** 2003. "Competition and Service Quality in the US Airline Industry." *Review of Industrial Organization*, 22(4): 275–96.
- Monteverde, Kirk, and David J. Teece.** 1982. "Supplier Switching Costs and Vertical Integration in the Automobile Industry." *Bell Journal of Economics*, 13(1): 206–13.
- Nickerson, Jackson A., and Brian S. Silverman.** 2003. "Why Aren't All Truck Drivers Owner-Operators? Asset Ownership and the Employment Relation in Interstate For-Hire Trucking." *Journal of Economics and Management Strategy*, 12(1): 91–118.
- Tadelis, Steven.** 2002. "Complexity, Flexibility, and the Make-or-Buy Decision." *American Economic Review*, 92(2): 433–37.
- Williamson, Oliver.** 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
- Williamson, Oliver.** 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. New York: Free Press.