

Control Rights, Network Structure and Vertical Integration: Evidence from Regional Airlines

Silke Januszewski Forbes
University of California, San Diego
sjanuszewski@ucsd.edu

Mara Lederman
Rotman School of Management, University of Toronto
mara.lederman@rotman.utoronto.ca

First Draft: December 2004

This Draft: February 2006

Abstract

This paper investigates the relationship between vertical integration and the importance of control rights under incomplete contracts. Our setting is the U.S regional airline industry. Regional airlines operate flights for major carriers under the major's brand. The majors market the regionals' flights as their own. There is substantial heterogeneity in whether or not regionals are owned by the major for which they operate. Furthermore, several majors own some of their regional partners while also contracting with others. We develop a simple framework that illustrates the benefits and costs of vertical integration between a major and regional. We argue that when unforeseen disruptions create the need for schedule adjustments – as frequently occurs in the airline industry - the major will internalize the impact of the disruption on its entire network, while the regional will not. Ownership of a regional mitigates this incentive problem by giving the major rights of control over how the regional's physical assets and labor force are used. However, by bringing the regional's labor force "in-house", ownership of a regional may erode some of the labor cost savings that are very reason why majors subcontract certain flights to regionals. Using data on majors' use of regionals in the second quarter of 2000, we test whether majors' choice of organizational form reflects this tradeoff between greater control and lower labor costs. Our results provide support for our analytical framework. We find that (1) owned regionals are more likely to serve city pairs that are more integrated into the major's network, where externalities not internalized by the regional will be the greatest; and (2) owned regionals are more likely to serve city pairs with adverse weather conditions, where unforeseen schedule disruptions will be more common.

*We thank Daniel Akerberg, Pierre Azoulay, George Baker, Vincent Crawford, Kenneth Corts, Martin Hellwig, Ignatius Horstmann, Thomas Hubbard, Darin Lee, Scott Masten, Wallace Mullin, Nancy Rose, Amedeo Odoni, Brian Silverman, Joel Sobel, Andrew Sweeting, Steven Tadelis, and seminar participants at the Massachusetts Institute of Technology, the University of California, Los Angeles, the University of California, San Diego, the UBC Sauder School of Business, the Rotman School of Management/Institute for Policy Analysis Conference on Organizational Economics, and the 2005 International Industrial Organization Conference for helpful comments and conversations. All errors are our own.

I. Introduction

Beginning with Coase's (1937) article, economists have been interested in understanding what determines the "boundaries of the firm". That is, how do firms decide which transactions to carry out in-house and which to procure through the market? In response to this question, a theoretical literature has developed which emphasizes the role of ownership in providing optimal incentives in an environment in which contracts are incomplete.¹ If the contracts that govern a firm's transactions in the market are incomplete, then unforeseen contingencies may be resolved in a manner which is not in the firm's best interest. By giving the firm ownership of the assets used in a transaction, vertical integration can mitigate this incentive problem since it gives the firm the right to decide how those assets are used in unforeseen circumstances.

In this paper, we study the relationship between vertical integration and the benefit of having residual rights of control when contracts are incomplete. We investigate this question in the context of the U.S regional airline industry. Regional airlines operate flights on short- and medium-haul routes for major carriers under the brands and codes of the majors. The majors ticket and market the regionals' flights as their own. There is substantial heterogeneity in whether or not regionals are owned by the major for which they operate. Furthermore, several majors own some of their regional partners while also contracting with others. This allows us to investigate the factors affecting the "make-or-buy" decision for a *given* firm. We show that the choice to vertically integrate is explained by market characteristics that affect the magnitude of the incentive problems that arise under incomplete contracts. In contrast to much of the previous empirical literature on this question, our empirical tests are not based on asset specificity or specific investments.² Rather, we find that vertical integration in our setting is influenced by the frequency of adaptation decisions and the existence of externalities across transactions resulting from the integration of transactions into a network.

¹ See, for example, Williamson (1971, 1985), Grossman and Hart (1986) and Hart and Moore (1990). Gibbons (2004) reviews some of the work that has followed these earlier contributions.

² Empirical studies that test the relationship between asset specificity and vertical integration include Monteverde and Teece (1982), Anderson and Schmittlein (1984), Masten (1984), Masten and Crocker (1985), Joskow (1985, 1987), and Hubbard (2001).

We begin by developing a simple framework that analyzes the benefits and costs of vertical integration in this industry. Our framework illustrates how these benefits and costs result from operational and institutional characteristics of the industry. Specifically, majors interact with their regionals during two types of operational decisions - *ex ante* scheduling decisions and real-time adjustments to schedule disruptions. Contracts between majors and regionals generally cover the first set of decisions; however, they do not - and likely could not - cover the second set of decisions. Thus, the primary incentive problem between majors and regionals results from this incompleteness of contracts with respect to real-time schedule adjustments. Unanticipated schedule disruptions are extremely common in this industry, resulting, for example, from adverse weather or mechanical problems. When these types of disruptions create the need for adjustments to the major's planned flight schedule, the major and its regional may disagree on what adjustments should be made. In particular, while the major will attempt to internalize the impact of the disruption on its entire network, the regional, who is compensated only based on the routes it serves for the major, will not. Ownership of a regional mitigates this incentive problem by giving the major residual rights of control over how the regional's physical assets and labor force are used. This, we argue, is the benefit of vertical integration in this industry.

However, there are also costs associated with vertical integration with a regional. Majors subcontract service to regional airlines because regionals have a cost advantage that results primarily from the lower salaries paid to regional airline employees, relative to the major's own employees. Ownership of a regional has the potential to erode this labor cost savings that regionals afford majors. The lower salaries paid to regional airline pilots have led these pilots to seek compensation that is closer to that earned by their counterparts at the mainline. Regional pilots' demands for higher wages may be harder for management to resist when a regional is wholly-owned by its major.

Our framework predicts that a major's optimal choice of organizational form will reflect the tradeoff between its incentive to exercise control over its regional and its incentive to maximize the labor cost savings that its regional provides. To test this framework, we develop two propositions that relate an airline's likelihood of using an owned regional on a city pair to airline-specific characteristics of that city pair which proxy for

the magnitude of the incentive problem. Our first proposition relates to the extent to which a regional's flight is integrated into the major's network. The more integrated a regional's flight, the more likely it is to experience disruptions and the more costly it will be for the major to have these disruptions resolved by its regional who will not internalize the impact of its decision on the major's network. Our second proposition relates to the frequency of unforeseen schedule disruptions that result from adverse weather. Adverse weather affects flight schedules by increasing the amount of time that is needed in between consecutive takeoffs or landings, thus forcing airlines to delay or cancel flights. As a result, adverse weather forces majors and regionals to make more frequent adaptation decisions. We therefore expect that wholly-owned regionals are more likely to be used on city pairs that are more integrated into the major's network and on city pairs that are more likely to be affected by adverse weather. We test these two propositions using a simple logit model. We use data on all flights served by regional carriers operating for the seven largest U.S network carriers on city pairs between the 300 largest U.S. airports in the spring of 2000. Since several of the major airlines use independent as well as vertically integrated regional carriers, we are able to estimate our model with airline fixed effects, using a subset of the data.

Our results suggest that vertical integration is used in this industry to mitigate incentive problems resulting from incomplete contracts. Furthermore, we find similar effects with and without airline fixed effects in the estimation, which suggests that our effects are identified by variation *within* firms as well as variation across firms. More specifically, we find that vertically integrated regional partners are more likely to be used on city pairs which have the major's hub at least at one endpoint. In addition, we find that they are more likely to be used on city pairs between airports at which the major operates a larger number of flights. We find this to be true even when we restrict the sample to routes that do not have the major's hub on either endpoint, indicating that this variable is not simply capturing the hub effect. With respect to our second proposition, we find that city pairs between airports with less favorable weather conditions tend to be served by owned regionals. These results are robust to several alternate measures of endpoint weather conditions. Specifically, we find that vertically integrated regionals are more

likely to be used on routes between airports with more rain overall and, in particular, with more rain during cold weather months.

This paper contributes to a growing empirical literature that seeks to explain patterns of vertical integration within an industry. It is perhaps most closely related to Baker and Hubbard (2004) which also provides evidence that allocating residual control rights through asset ownership can mitigate incentive misalignments that arise under incomplete contracts. It is also related to Nickerson and Silverman (2003) who explore the effects of externalities that arise in less-than-truckload carriage on the choice of organizational form. These externalities are quite similar to those that result here from airlines' hub-and-spoke systems. In addition, this paper is related to empirical work in the franchising literature (indeed, the relationship between a major and its regional is very similar to a franchise relationship) and to empirical work in the transaction-cost economics literature which investigates the role of contractual incompleteness by examining the relationship between asset specificity and vertical integration.³

We believe that this paper extends this literature in several important ways. First, from an empirical perspective, our setting has the benefit of allowing us to observe the same firm using alternate organizational forms for different transactions. This allows us to investigate the relationship between vertical integration and underlying market characteristics while controlling for unobservable factors that affect a particular firm's relative returns from vertical integration. To our knowledge, we are one of the first papers outside of the franchising literature to exploit this.

Second, our paper illustrates the close link between the choice of organizational form and the operational and institutional characteristics of an industry. In our setting, the incentive problem between the major and the regional results primarily from the network structure of the industry. Every decision made concerning one part of the network has implications for other parts of the network. The challenge associated with subcontracting a portion of the network - as majors do with regionals - is that when unforeseen contingencies arise, the subcontractor will, in general, not have the incentive to internalize the impact of its actions on the remainder of the network. We suspect similar incen-

³ Lafontaine (1992) and Lafontaine and Slade (1997) study the role of agency costs for vertical integration in franchising. Empirical studies on asset specificity and vertical integration are referenced in footnote 2.

tive problems may arise when firms attempt to subcontract in other network-based industries.

Finally, while the benefits of vertical integration result from the network structure of operations in this industry, the costs of vertical integration in our setting result from the nature of labor relations in this industry. Conversations with industry participants revealed that majors were very much concerned that ownership of a regional would jeopardize the pay differential that exists between major and regional pilots. Thus, we believe our setting suggests the importance of considering not only the “standard” incentive costs of vertical integration but also other costs which may result from the unique institutional characteristics and history of an industry.

The remainder of this paper is organized as follows. The next section provides the institutional details that inform our analysis of the costs and benefits of vertical integration in this industry. Section III presents the analytic framework and develops our empirical propositions. In Section IV, we describe the data and variables. In Section V, we explain the empirical approach and present the results. A final section concludes.

II. Organizational Forms in the Regional Airline Industry

In this section, we provide the institutional details that inform our analysis of the costs and benefits of vertical integration in this industry. We begin by describing the role of regional airlines as “subcontractors” for major U.S. network carriers. We then describe the two organizational forms that govern relationships between majors and regionals - contracts and ownership - and highlight important differences between the two. We conclude with a brief discussion of the contracts used by majors and independent regionals.

II.A. The Role of Regional Airlines

Regional airlines operate as “subcontractors” for major U.S. network carriers on short and medium-haul routes. These routes are typically low-density and most efficiently served by small aircraft. Almost all regional airlines operate under codeshare

agreements with one or more major carriers.⁴ Under these agreements, the regional operates flights on behalf of the major carrier, who 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, the regional's planes are painted in the major's color scheme, passengers traveling on the regional earn the major's frequent flyer points, and the regional uses the logos, trademarks and even the name of the major (for example, Delta's regional Comair operates under the name Delta Connection).⁵ To facilitate passenger connections between a major and its regional, their schedules, as well as check-in and baggage handling, are typically coordinated.

In this capacity, regional airlines have come to play a vital role in the U.S. commercial airline industry, providing the sole means of scheduled air transportation at more than two-thirds of all North American airports served by U.S. carriers. Table 1 provides some descriptive data about the activity of regional airlines in 2000 (the year of our sample). In 2000, there were 94 regional airlines in operation, though 99% of activity was accounted for by the 50 largest regionals. Regionals served a total of 729 airports in North America and completed 4.46 million departures. The average passenger trip on a regional was about 300 miles and the average seating capacity of aircraft operated by regionals was 31.7.

Majors subcontract service to regional airlines because regionals have a cost advantage on the types of routes that they serve. This cost advantage results primarily from the lower salaries paid to regional airline employees, relative to the major's own employees.⁶ Regional airlines' lower labor costs can be traced to their origins as non-unionized and non-regulated operators of small aircraft. Like majors, many regionals now have unionized workforces; nonetheless, this labor cost differential has persisted.⁷ In addition to

⁴ In 2003, 99% of regional airline passengers traveled on flights that were codeshared with a major carrier.

⁵ Note that this is different from the type of codeshare arrangement typically negotiated between two major carriers, such as United Airlines and Lufthansa. Under that type of agreement, both carriers will sell tickets under their own codes on each other's flights but the operating carrier will maintain its own identity.

⁶ Salaries are not directly comparable because major airlines fly larger aircraft than regional carriers, but hourly pilot salaries for the smallest equipment flown by major airlines are about twice as high as hourly pilot salaries for the largest equipment flown by regional carriers, controlling for the years of experience that the pilot has.

⁷ With the introduction of the regional jet in the late 1990s, the planes flown by regionals are now much more similar to those flown by the mainline and the case for the pay differential is arguably weaker. Pilots at the mainline recognize this increased substitutability between regionals' flights and their own and nego-

the advantage provided by regionals' lower labor costs, there may be gains (in the form of lower maintenance and training costs) to having the major and the regional each specialize their fleet to include only a small number of different aircraft types. In fact, both owned and independent regionals perform their own aircraft maintenance rather than have it done by the majors.

The role of regional airlines as cost-effective providers of air service to small communities can be traced back to the pre-deregulation era.⁸ Beginning in the late 1960s, small commuter airlines started offering service to small communities which the "trunk" and "local service" carriers had previously been serving with large government subsidization. Deregulation brought a reorganization of the large airlines' networks from point-to-point to hub-and-spoke. The large airlines quickly realized that passengers traveling to or from small communities on commuter airlines tended to fly these trips as part of longer itineraries and, as such, could provide an additional source of feeder traffic at their hubs. Coordination with the commuter leg of these passengers' itineraries would help them capture these passengers for the second leg of their trip. This presented the large airlines with two options: serve the short-haul routes themselves or establish arrangements with the existing commuter carriers. Even with the improved efficiency brought by deregulation, the larger airlines' costs of serving these small towns were still well above those of the specialized commuter airlines. As a result, partnerships between majors and commuter airlines (now called regional airlines) emerged.

II.B. Organizational Forms

Today's 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. In the case of a wholly-owned regional, "vertical integration" means that the major carrier owns the assets of the regional but the re-

tiate so-called "scope clauses" to restrict the number and types of routes which may be subcontracted to regional carriers.

⁸ See Borenstein and Rose (2005) for a thorough account of regulatory reform in the U.S. airline industry. See Levine (1987) pages 437-441 for a discussion of the evolution of partnerships between major airlines and commuters airlines.

gional and the major maintain separate operations and labor contracts.⁹ Note that wholly-owned regionals could, in principle, also contract with major carriers other than the one by which they are owned, but we never observe such relationships in our data.

Table 2 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. Some majors own all of their regional partners, others own none and yet others use a mix of owned and independent regional carriers. In 2000, American Airlines was the only major network carrier to use only wholly-owned regionals. American owned both American Eagle and Business Express, both of which operated under the American Eagle brand. In contrast, United Airlines and TWA both used only independent regionals. United Airlines, for example, contracted with four independent regionals, each of which operated under the United Express brand. Continental, Delta, Northwest and US Airways are perhaps the most interesting because they each used a mixture of owned and independent regionals. Delta, for example, had contract relationships with three independent regionals but also had two wholly-owned regionals, all of which operated under the Delta Connection brand.

II.C. Differences between Owned and Independent Regionals

While both owned and independent regionals operate as subcontractors for majors, there are a number of important differences between the two. These differences, combined with assumptions about contractual incompleteness, generate the costs and benefits of vertical integration discussed in the next section.

⁹ Separate operations are necessary so that the major can legally maintain distinct labor contracts (one for its own employees and one for each regional's employees) and thereby preserve the cost advantages that regionals have. If two separate airlines are effectively being operated as a single entity, the unions representing employees at those airlines may file an application with the National Mediation Board (NMB) seeking to have them declared a "single transportation system". If their application is granted, the unions of the carriers will operate as a single entity. To our knowledge, there have been no such applications by the unions of a major and its regional, however industry participants indicate that the risk of a "single transportation system" designation is the reason why majors operate owned regionals as separate operations.

First, independent regionals own or lease their own aircraft and hire, fire and manage their own employees. Wholly-owned regionals - though operating as a separate entity within the major - ultimately have their aircraft and employees included as part of the major's own fleet and workforce. As such, independent regionals retain residual rights of control over their aircraft and workforce while, for wholly-owned regionals, these rights ultimately rest with the major. In addition, ownership of a regional carrier allows the major airline to select and replace the regional's management while contracting with an independent regional does not. The implication of this is that the managers of a wholly-owned regional are ultimately accountable to the major.

Second, ownership affects the way in which majors and regionals respond to unanticipated schedule disruptions.¹⁰ These disruptions occur most frequently in adverse weather conditions, but may also result from air-traffic control problems or airline mechanical problems. When an airport experiences adverse weather, the Federal Aviation Administration (FAA) will determine - several hours in advance - the number of flights which will be allowed to land during each hour. When weather necessitates a reduction in flights relative to the original schedule, each airline will receive a number of take-off and landing slots in proportion to its original share of scheduled flights. The airline then decides which of its flights to delay or cancel. When a regional is owned by a major, the major and the regional receive a *common* allocation of slots and the major carrier decides which of its own and which of the regional's flights to delay or cancel. In fact, the re-scheduling decisions for wholly-owned regionals are done by the major carrier's Airline Operational Control Center (AOCC). In contrast, when a regional is independent, it receives its own slot allocation and makes its own decision (in its own AOCC) about delays and cancellations, possibly in coordination with the major carrier for which it operates.

Third, there may be operating cost differences between wholly-owned and independent regionals.¹¹ The lower salaries paid to regional airline pilots have led these pilots to seek compensation that is closer to that earned by their counterparts at the mainline. Regional pilots' demands for higher wages may be harder for management to

¹⁰ Ball *et al.* (2005) discuss airport and airline responses to schedule disruptions in detail.

¹¹ Both types still have a substantial cost advantage over majors.

resist when a regional is wholly-owned by a major airline.¹² Indeed, the case for a common labor contract for both sets of pilots is stronger when the two sets of pilots are part of the same organization. In addition to the difference in wages, phone conversations with industry executives and analysts have suggested that owning a regional may lead to costs associated with managing two distinct labor forces, such as more frequent labor disputes.

II.D. The Nature of Contracts

In Section III, we develop a simple framework of the costs and benefits of vertical integration with a regional partner. Because that analysis depends importantly on the assumption that contracts in this industry are incomplete, we will briefly describe what we know about contracts in this industry.¹³

Contracts between majors and independently owned regionals generally take one of two forms. First, contracts may be structured as revenue-sharing agreements, under which the regional carrier agrees to serve a set of routes on behalf of the major and to coordinate its schedule on those routes with the major's own schedule. Under these agreements, the fares are set by the major carrier, and the regional receives an allocated portion of the revenue from passengers traveling a portion of their itinerary on the regional.¹⁴

Alternatively, contracts may be fixed-fee agreements, also known as capacity-purchase agreements. Under these contracts, 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. This fixed payment is calculated to cover the regional's operating costs and to guarantee a reasonable rate of profit. In addition, the regional may receive incentive payments based on operational performance, such as on-time performance and baggage handling. Compared to revenue-sharing,

¹² For example, after Delta acquired the previously independent regional carrier Comair, pilots at Comair demanded higher salaries and went on strike. The pilots used the now common ownership of Delta and Comair as the main argument for a salary increase.

¹³ We do not have access to actual contracts between majors and regionals but have based the description that follows on information collected from the following sources: conversations with industry participants, trade press articles, annual reports of majors and publicly traded independent regionals, and documents relating to a breach of contract lawsuit between United Airlines and Mesa Airlines in 1997.

¹⁴ The regional receives 100% of the revenue from passengers traveling entirely on the regional (i.e.: passengers who do not connect to the major's flights).

fixed-fee contracts reduce the risk that the regional faces, but also reduce its incentives to lower costs and increase demand. From the major's perspective, these contracts provide it with greater control over its regional, in particular over its schedule, since the regional's willingness to operate a particular flight is independent of cost and demand conditions.

Note that both types of contracts were in use during our sample period. We do not, however, have systematic data on which independent regionals were covered by a particular contract form. Instead, we assume that majors choose to use the contract that is optimal. While fixed-fee contracts may improve upon certain incentive problems that arise as under revenue-sharing contracts, as we explain below, our analysis of the costs and benefits of vertical integration do not depend on which type of contract was in place.

III. Theoretical Considerations

In this section, we analyze the benefits and costs of owning a regional. We begin by describing the two types of operational decisions that airlines make and the extent to which each type can be contracted on. We then explain how the inability of majors and regionals to contract on certain types of operational decisions creates an incentive problem that is mitigated by a major's ownership of its regional. This, we argue, is the benefit of vertical integration in this industry. A major's optimal organizational form will reflect a tradeoff between this benefit and the higher labor costs that may result from ownership of a regional. We conclude the section by developing our empirical propositions.

III.A. Operational Decisions

Before proceeding, it is useful to distinguish two different types of decisions that airlines make. The first set of decisions is airlines' *ex ante* scheduling decisions. Airlines determine their schedules in advance based on expected demand, expected costs and their expectations of competitors' decisions. These scheduling decisions generate a set of operational decisions such as aircraft and crew allocation, and fuel and catering orders. Majors interact with their regional during these *ex ante* scheduling decisions when they decide which flights to subcontract to the regional for operation.

The second set of decisions is airlines' real-time adjustments to their planned schedules. Real-time adjustments to airlines' planned schedules are common and can arise for a large number of reasons. For example, they may arise from mechanical problems, adverse weather, or security or air traffic control disruptions. The network structure of airlines' operations means that disruptions in one part of an airline's network may require the airline to make real-time adjustments at other points of its network. Furthermore, the network structure of airlines' operations means that any real-time adjustments that an airline makes must be based on a re-optimization of the airline's entire network, not just a re-optimization of the flight or route in question. For example, the decision to cancel a flight in response to bad weather must consider not only the costs of canceling that flight, but also the costs associated with the aircraft and crew from that flight not being available at the arrival airport for their next scheduled departure. The re-optimization of aircraft, crew, and passenger schedules is a difficult operations research problem for which, currently, only heuristic solutions are available.¹⁵ Majors interact with their regionals during these types of decisions when disruptions require majors to make real-time adjustments to their flights which are operated by the regional

III. B. Incentive Problems between Majors and Regionals

Contracts between majors and regionals will generally cover the first type of decisions but not the second. That is, majors and regionals do contract on *ex ante* scheduling decisions; however, they do not (and likely could not) contract on real-time schedule adjustments. The full set of possible schedule disruptions is so large (and difficult to anticipate) that it would likely not be possible for the major and the regional to specify a contract covering all possible contingencies.¹⁶ Moreover, because airlines' response to disruptions require a re-optimization of their network, contracting on these contingencies may even be undesirable since it would eliminate valuable flexibility.¹⁷

¹⁵ See Ball *et al.* Furthermore, since there are large fixed costs of developing solution algorithms, the major carriers have developed better algorithms than smaller airlines.

¹⁶ Williamson (1985) argues that bounded rationality is an important reason why contracts are incomplete. Williamson (1971) and Grossman and Hart (1986) point out that when contracts are incomplete, common ownership of assets – or vertical integration – creates incentives that cannot be replicated through contracts.

¹⁷ Williamson (1975) argues that vertical integration can maintain flexibility which would be eliminated by formal contracts. Tadelis (2002) emphasizes that it may be desirable not to specify complex decisions *ex*

Thus, the primary incentive problem between majors and regionals results from this incompleteness of contracts with respect to real-time schedule adjustments. When disruptions create the need for adjustments to the major's planned flight schedule, the major and its regional may disagree on what adjustments should be made. More specifically, the major's optimal response to a disruption will attempt to internalize the impact of the disruption on its entire network while the regional's (who is compensated only based on the routes it serves for the major) optimal response will not. For example, adverse weather at an airport may require an airline to cancel some of its flights to that airport. A major's decision of which flights to cancel will consider the number of passengers onboard, the number of those making connections, the planned routing of the aircraft and crew, and any effects of the cancellation on future demand. However, its regional - who is compensated only based on the routes it serves - will be only be concerned with maximizing route-level profits (if covered by a revenue-sharing agreement) or maximizing completion factors and on-time performance (if covered by a fixed-fee contract). Similar incentive problems may arise when majors have to make real-time decisions about delaying outgoing flights in response to late incoming flights, combining flights in response to low load-factors, substituting aircraft in response to mechanical problems or prioritizing gate and ground crew access.

When majors contract with independent regionals, these types of unforeseen schedule disruptions may be resolved in ways that are not optimal for the major. Alternatively, resolution of these types of disruptions may require costly renegotiations between the major and its regional.¹⁸ Ownership of a regional mitigates these incentive problems by giving the major residual rights of control over how the regional's physical assets and labor force are used.¹⁹ This is the benefit of vertical integration in this industry.²⁰

ante in a contract in order to preserve flexibility and reduce the costs of negotiations over *ex post* adaptations.

¹⁸ Recall that independent regional have their own AOCCs to carry out schedule re-optimizations while owned regionals' rescheduling is done by the major's AOCC.

¹⁹ The link between asset ownership and residual rights of control plays a central role in Grossman and Hart's theory. Baker and Hubbard (2004) show that asset ownership is used in the trucking industry to mitigate incentive problems that arise from incomplete contracts.

²⁰ Note that in addition to incentive misalignments with respect to real-time schedule adjustments, vertical integration may solve a second incentive problem between regionals and majors. This incentive problem results from the fact that an independent regional will not internalize the effects of its actions on the value of the major's brand. Thus, if monitoring is imperfect, the regional will provide a suboptimal level of qual-

The costs of vertical integration in this industry were highlighted in Section II. Ownership of a regional has the potential to erode the labor cost savings that regionals afford majors. Preservation of the labor cost differential requires the major to manage its own employees and the regional's employees under two distinct labor contracts. Ownership of the regional by the major may reduce the distinction between the two sets of employees, both in the eyes of the regional pilots and in the eyes of the National Mediation Board.

III. C. Empirical Propositions

A major's optimal choice of organizational form will reflect the tradeoff between its incentive to exercise control over its regional and its incentive to maximize the labor cost savings that its regional provides. The cost disadvantage of an owned regional should be constant across routes, after controlling for distance, because pilots are paid based on block hours. If ownership mitigates the incentive problem that exists when a major contracts with a regional, then we should observe that the likelihood of using an owned regional is increasing in the size of the incentive problem for a given city pair. On the other hand, if formal contracts on their own - or in combination with relational ones - can mitigate the incentive problem, we should find no effect of variables that measure the size of the incentive problem on the likelihood of ownership. We develop two propositions that relate an airline's likelihood of using an owned regional on a city pair to airline-specific characteristics of that city pair which proxy for the magnitude of the incentive problem. We assume that the incentive problem is greater (and, hence, the value of control that comes with ownership larger) when either (a) the likelihood of unanticipated schedule disruptions is higher; or (b) the costs to the major of having real-time adjustments made in a suboptimal way are greater.

Our first proposition relates to the extent to which a regional's flights are integrated into the major's network. Flights that are more integrated into the major's network are scheduled to facilitate passenger connections between those flights and other flights operated by the major. We expect the incentive problem between a major and its

ity. Note that this problem is parallel to that which exists between franchisors and franchisees. We do not focus on this benefit of vertical integration because we have no variation across city pairs in the size of this incentive problem.

regional to be larger on these types of flights for two reasons. First, the more integrated a flight is into the major's network, the more likely that flight is to be affected by unanticipated disruptions. This is because the flight can be delayed or cancelled not only for reasons that concern it directly, but also because of disruptions that occur elsewhere in the major's network. Second, the more integrated a flight is into the major's network, the more costly it will be for the major to have disruptions relating to that flight resolved by its regional, when the regional does not internalize the externality of its decision on other flights in the majors network. For these two reasons, we propose that *city pairs that are more integrated into the major's network are more likely to be served by a wholly-owned regional.*

Our second proposition relates to the frequency of unanticipated schedule adjustments that result from a particular cause - adverse weather conditions (such as rain, fog or snow). Weather is the leading cause of flight delays that is not under the airline's control. Adverse weather affects flight schedules by increasing the amount of time that is needed in between consecutive takeoffs or landings. Flight schedules are set assuming weather conditions that are favorable for flying. When the amount of time needed between flights increases, takeoff and landing slots are reduced and airlines are forced to delay or cancel flights. This makes unanticipated schedule adjustments necessary. For these reasons, we propose that *city pairs that are more likely to experience adverse weather are more likely to be served by a wholly-owned regional.*

IV. Data

IV.A. Data and Sample

The empirical analysis is based on airline schedule data compiled by the Official Airline Guide (OAG). These data provide the complete weekly flight schedules of all domestic airlines. A representative week is provided for each quarter. Each observation in the data corresponds to a particular flight by an airline in a quarter (for example, American Airlines flight #596 between Chicago O'Hare and Atlanta in the first quarter of 2000). For each flight, the data provide information on the carrier, the origin and destination airports, the scheduled departure and arrival times, the days of operation, the type of aircraft used, and, for flights that are operated by regional partners, the identity of the re-

gional carrier. Recall that as described above, for flights operated by regional partners, all tickets are sold under the name of the major.

We supplement the OAG data with two additional data sources. First, we use data from the Regional Airline Association (RAA) to determine, for each flight that is operated by a regional, whether that regional is owned by the major or independent. The RAA provides annual information on every major domestic airline's regional partners in that year and whether or not that partner is wholly owned by the airline for which it is operating. Second, data on average monthly temperatures and precipitation levels at each airport location are taken from the Spatial Climate Analysis Service.

Our sample period is the second quarter of 2000. We choose this quarter of the year because it contains the fewest number of seasonal flights. We consider all flights flown by regionals between the top 300 U.S. airports.²¹ Because we are interested in analyzing airlines' form of organization on a particular flight segment, our level of observation is the airline-city pair. We use the term "city pair" to refer to direct non-stop service between two endpoint airports, in either direction (i.e.: a flight from LAX-ORD and a flight from ORD-LAX are considered the same city pair). 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. On these routes, majors will use regionals to operate the flights at off-peak times of day when demand is expected to be lower in order to maintain high flight frequencies on the route. For city pairs that are served by both the major carrier and a regional partner, we investigate the type of regional partner that is used for the subset of flights operated by the regional. Our empirical results are robust to excluding city pairs on which the major operates some fraction of flights.

We consider the organizational form decisions of the seven largest network carriers: American Airlines, Continental Airlines, Delta Air Lines, Northwest Airlines, Trans World Airlines, United Airlines and US Airways. We focus on the largest network carriers because many of the predictions from the theory relate ownership decisions to charac-

²¹ 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 either endpoint in Alaska, Hawaii, Puerto Rico, Guam or the U.S. Virgin Islands.

teristics of airlines' hub operations. The results are robust to including smaller network carriers such as Midway Airlines and Midwest Express. We analyze flights operated by regional carriers for these major airlines. 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.²² We restrict our analysis to flights that operate on Mondays, thus only considering the airlines' weekday schedules.²³ These data restrictions result in a sample of 994 airline-city pairs.

IV.B. Variables

This section describes our variable constructions. Variable names and definitions appear in Table 3. Summary statistics appear in Table 4.

Our dependent variable is OWNED_REGIONAL, a dummy variable which equals one if a flight is operated by an owned regional partner. To test our first proposition, we construct two different measures of the extent to which a flight is integrated into the major's network. Our first measure is the dummy variable HUB, which equals one if either endpoint of the city pair is the major's hub. On airline-city pairs which involve a hub, 60 percent of all passengers connect to or have connected from another flight, compared 26 percent on airline-city pairs which do not involve a hub.²⁴ We expect that majors will be more likely to use owned regionals on city pairs that have a hub on at least one endpoint. Although, as described in Section II, regionals are primarily used to provide feeder traffic to majors' hubs, 30% 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.

To further investigate the relationship between organizational form and the integration of a city pair into an airline's network, we measure the total number of flights that the major operates from the endpoint airports of a city pair on a given day. One advan-

²² For example, we exclude the limited codeshare relationship between Northwest and American Eagle in which Northwest sells tickets on some American Eagle flights out of Los Angeles but these flights are operated under the American Eagle brand in the sense that the plane exterior and other marketing materials are associated with American Eagle, not Northwest.

²³ 99.46% of the flights which operate on Mondays also operate Tuesday through Friday.

²⁴ These estimates are based on passenger numbers from the Department of Transportation DB1A database for airline-city pairs in our sample.

tage of these variables is that they allow us to investigate the relationship between an airline's network at an airport and organizational form using a subsample that includes only non-hub airports. They also provide a continuous measure of network size. We construct `LARGER_FLIGHTS` which measures the number of flights the major operates out of the endpoint at which it is larger (i.e.: operates more flights) and `SMALLER_FLIGHTS` which measures the number of flights the major operates out of the endpoint at which it is smaller. Both are measured in hundreds of flights and do not include flights operated by the major's regionals.

To test our proposition relating to adverse weather, we construct three weather variables that measure expected weather conditions on a particular city pair²⁵ First, for each endpoint airport, we calculate the average number of months per year in which the average daily minimum temperature at that airport is below zero degrees Celsius. We then take the maximum of this variable across the two endpoint airports of a city pair and call this `FREEZING`. Second, for each endpoint airport, we calculate the average annual precipitation at that airport. We again take the maximum of this variable across the two endpoint airports of a city pair and call this `PRECIP`. Finally, for each endpoint airport, we calculate the average annual precipitation during months in which the average daily minimum temperature is below zero degrees Celsius. We take the maximum of this across the two endpoint airports and call this `FREEZ_PRECIP`. This third variable captures the fact that it is largely the coincidence of cold temperatures and high precipitation that causes schedule disruptions.

IV.C. Descriptive Analysis

Table 4 contains summary statistics for the variables described above. The summary statistics indicate that slightly over half of the flights served by a regional are served by one that is wholly owned by its network carrier partner.

Analyzing the decision whether to use a regional partner (of any type) on a route is not the focus of this paper. However, to provide some intuition for the types of routes majors subcontract to regionals, Table 5A compares the characteristics of routes less than

²⁵ As we discuss below, we also construct alternate weather measures and use these to test the robustness of the results.

1500 miles²⁶ served by majors with those served with either type of regional partner. For this purpose, we include two additional variables in the table, the distance of the route measured in miles and the arithmetic mean of the populations of the endpoint cities which we call MEAN_POP. The population variable is constructed from the 2000 Census of Population. The first thing to note from this table is that – even in this restricted set of city pairs – routes served by regionals are much shorter (on average 314 miles) than those served by majors themselves (on average 730 miles). City pairs served by majors themselves are more likely to have a hub on at least one endpoint and are more likely to be between larger airports (as measured by the major’s number of other flights at the airport). As well, city pairs served by majors tend to be between cities with larger populations. There are no large differences in the average weather patterns of city pairs served by majors and those served by regionals. Overall, the descriptive statistics suggest that a major’s choice between serving a city pair itself or with a regional partner is largely driven by route density and distance.

Table 5B compares the characteristics of city pairs served by owned and independent regionals. Examination of the unconditional means in this table provides some preliminary evidence on the hypotheses laid out above. The means suggest that city pairs served by owned regionals are more likely to have the major’s hub on at least one endpoint. In addition, city pairs served by owned regionals are between endpoint airports at which the major operates a larger number of flights. Finally, city pairs served by owned regionals have greater precipitation at the endpoint airports. Note that there are only small differences between the city pairs served by the two types of regionals with respect to endpoint population and distance. This suggests that these variables determine the choice between the use of the major’s own planes and that of a regional partner but do not influence the choice of the type of regional that is used.

V. Estimation and Results

V.A. Empirical Approach

We investigate a major airline’s decision to use an owned rather than independent regional partner on a particular city pair, conditional on the major using a regional

²⁶ This is the maximum range for planes flown by regional carriers.

carrier (of either type) to serve that city pair.²⁷ We relate a major's decision about what type of regional to use to city pair characteristics (some of which are airline specific) that proxy for the magnitude of the incentive problems described in Section III. In particular, we have two types of variables that capture variation in the magnitude of incentive problems across city pairs. These are the network variables and the weather variables. The parameters on these variables are identified by relating variation in these characteristics to both cross- and within-carrier variation in the type of regional used. In some specifications, we exploit the fact that we have within-airline variation in organizational form and include only the four major carriers that used both types of regional partners in 2000. In these specifications, we include a dummy variable for each major carrier to control for that carrier's unobserved propensity to use a given type of regional. The results from the larger sample are generally consistent with the results from this subset of airlines. All specifications are estimated using a simple logit model.²⁸ We control for the distance of the city pair in all specifications to capture the fact that the cost disadvantage of owned regionals may increase with distance because pilots are generally paid based on block hours flown. The key identifying assumption of our empirical approach is that none of our explanatory variables are correlated with other unobserved factors that affect the relative returns to using an owned or independent regional. 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.

V.B. Results

i. Route Level Estimation

We now describe the results of our empirical estimation. All tables present marginal effects. For dummy variables, the reported effects are for changing the value of the

²⁷ We have also estimated the model as a nested logit model in which the major first decides whether to serve a city pair itself or with a regional and then, conditional on choosing a regional, decides whether to use an owned or independent one. Consistent with the means in Table 5A, we find that the decision whether to use a regional is largely determined by the distance and endpoint populations of the city pair. The results from that model on the choice between owned and independent regionals are consistent with the results presented here.

²⁸ The results are consistent th those obtained using a linear probability model.

explanatory variable from 0 to 1. All regressions include a constant term, but we do not report the coefficients on the constant. Table 6 presents our first set of results. The dependent variable is OWNED_REGIONAL, a dummy variable that equals one if the route is served by an owned regional and zero if the route is served by an independent regional. The sample includes all routes served by a regional of either type. Column (1) includes only the HUB variable, which tests our proposition that incentive problems are greater for routes that are more integrated into the major's network, and the distance of the route. 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. Having its hub on either endpoint increases a major's probability of using an owned regional by 19.9 percentage points. This finding is consistent with our hypothesis that control rights which the major acquires through vertical integration are more important on city pairs that are closely integrated into the major's network. The coefficient on distance is positive and significant at the 10% level in this specification, but it is not distinguishable from zero in any of the remaining specifications in this table. This reflects the fact that there is little variation in the distances of the routes served by owned and independent regionals.

Columns (2) and (3) add the weather variables. Recall that these measure the maximum of three weather characteristics across the two endpoint airports of a city pair. Column (2) includes only the direct effect of cold weather and precipitation but does not include FREEZ_PRECIP, which measures the total precipitation during cold months. The results show that the weather variables have highly significant effects on the type of regional chosen. Owned regionals are more likely to be used on city pairs involving airports with higher annual precipitation. At the sample mean, the estimate implies that an increase in annual precipitation of 1 mm would increase the likelihood of using an owned regional by 0.08 percentage points. An increase in annual precipitation by one standard deviation from the sample mean would increase the likelihood of using an owned regional by 22 percentage points. This effect is consistent with our hypothesis that vertically integrated regional carriers are more likely to be used on city pairs on which adaptation decisions have to be made more frequently. Cold temperatures - measured by the FREEZING variable - decrease the likelihood that the major uses an owned regional. The estimate implies that at the sample mean an additional month with average minimum

temperature below freezing would decrease the likelihood of using an owned regional by 4.09 percentage points. An increase in the number of cold months of one standard deviation from the sample mean would decrease the likelihood of using an owned regional by 7.0 percentage points.

While the finding that cold weather decreases the likelihood of using an owned regional may initially seem surprising, note that cold weather on its own does not prevent flying. To confirm that our weather variables are indeed capturing the likelihood of adaptation decisions due to weather-related schedule disruptions, we investigate the relationship between weather and actual delays. We describe the data and sample used for this analysis in more detail in Appendix A. The results, presented in Table A.1, indicate that, consistent with our findings in Table 6, actual flight delays *decrease* with the number of cold months at the airport and *increase*, as expected, with the amount of total annual precipitation and, especially, with precipitation during cold months. Thus, the relationship we estimate between weather and organizational form is consistent with the relationship we estimate in the Appendix between weather and actual delays. That is, to the extent that actual delays are a good proxy for the frequency of adaptation decisions, the results in Appendix A indicate that empirically airlines do make adaptation decisions *less* often at airports with more months of below freezing temperatures.

However, this begs the question of why flight delays are shorter in regions with colder weather. Once we control for the amount of precipitation during cold months, the effect of the number of months with below freezing temperatures on flight delays is identified by comparing warm and dry regions (generally the Southwest) to similarly dry but colder regions (such as the Mountain region, including states such as Colorado, Idaho, Montana, and Wyoming). A possible explanation for our finding that flight delays decrease with the number of cold months is that the population in the Southwest has grown very rapidly during the past decades, whereas the Mountain region has had less rapid population growth. Airport capacity, however, has not grown very much at all in either of these regions. We therefore suspect that the negative coefficient on the FREEZING variable is capturing the lower level of airport congestion in cold and dry areas.

Rather than cold temperatures alone, it is the coincidence of cold temperatures and precipitation (i.e.: snow and ice) that causes most flight delays and cancellations. To cap-

ture this effect, column (3) adds a control for the maximum (across the two endpoint airports of a city pair) amount of precipitation during cold months, FREEZ_PRECIP. The coefficient on this variable is positive and significant suggesting that while cold weather on its own does not increase the use of owned regionals, precipitation during cold months does. The coefficient on the number of cold months stays negative and increases in magnitude. The total annual precipitation is still estimated to have a positive effect.

In column (4), we add the other network variables, LARGER_FLIGHTS and SMALLER_FLIGHTS. Recall that the hypothesis about these variables is that they provide an additional measure of the integration of a city pair into the major's network. Column (4) shows that these variables have no statistically significant effect on the likelihood of vertical integration when the HUB variable is included in the regression. This, however, may be due to the high correlation between HUB and the number of flights, especially for the larger airport. Therefore, in column (5), we estimate the same specification excluding the hub variable. We now find that the number of flights at the larger endpoint has a positive effect on the likelihood of using an owned regional which is consistent with our hypothesis that incentives for integration are greater on city pairs that are more integrated into the major's network. The number of flights at the smaller airport, however, remains statistically insignificant. This suggests that the network integration matters at the larger endpoint, but not at the smaller endpoint.

Overall, we find evidence in favor of our hypotheses that city pairs which have characteristics that imply greater incentives for the major to obtain residual control rights are more likely to be served by a vertically integrated regional carrier. In order to control for unobserved factors that may have airline-specific effects on the relative returns from one type of regional over the other, we now restrict the sample to the major airlines which use both types of regionals and include a dummy variable for each of these major carriers. These airlines are Continental, Delta, Northwest, and US Airways. We estimate the same specifications as in Table 6 on this reduced sample. These results are reported in Table 7.

When we include dummies for all major airlines, the estimated effects of the network and weather variables have the same signs as in Table 6 and remain statistically significant. Looking at column (1) of Table 7, the hub effect is now estimated to be

about twice as large the one estimated in the same specification without airline fixed effects. The distance of the route has no statistically significant effect in any of the specifications in this table.

We find in column (2) that the number of cold months has a positive effect on the likelihood of owning a regional carrier. However, when we control for the amount of precipitation during cold months separately in column (3), the effect of FREEZING is negative as it is in all specifications of Table 6 and the remaining specifications in this table. We find that total annual precipitation and precipitation during cold months continue to have a positive effect on the likelihood of vertical integration. In fact, these effects are robust throughout our entire analysis. Notice also that the magnitude of the marginal effects is quite similar to the magnitudes estimated without airline-specific effects in Table 6.

Again, we add controls for the number of flights at each endpoint airport. In column (4), we find a marginally significant negative effect of the number of flights at the larger endpoint. When we drop the control for HUB in column (5), however, this effect is positive and highly significant, as before. The number of flights at the smaller endpoint remains insignificant in both of these specifications.

In Tables 8 and 9, we split the sample into hub and non-hub city pairs and re-estimate the specifications from Table 6, with the hub variable, of course, omitted. This allows us to investigate whether the other variables affect hub and non-hub city pairs in similar ways. However, because over 70% of the routes in our sample are hub city pairs, the sample size for the non-hub regressions is quite small. Looking first at hub city pairs (Table 8), the results are quite similar to those in Tables 6 and 7. The weather variables all have the same signs as before. The estimates from the specification in column (2), which includes FREEZING, PRECIP, FREEZ_PRECIP, and DISTANCE as explanatory variables, imply that at the sample mean an additional cold month decreases the likelihood of using an owned regional on hub routes by 13.2 percentage points, and an increase in annual precipitation by 1 mm would increase the likelihood of using an owned regional by 0.08 percentage points. An increase in the number of cold months by one standard deviation would decrease the likelihood of using an owned regional by 22.4 percentage points, whereas an increase in precipitation by one standard deviation would in-

crease the likelihood of using an owned regional by 23.09 percentage points. LARGER_FLIGHTS and SMALLER_FLIGHTS have no significant effect on the likelihood of using an owned regional on hub routes.

Table 9 presents the results for non-hub routes. Note that there are only 294 such airline-city pairs in our sample. Several of the estimates lose significance in some of the specifications, although the signs are consistent with those estimated earlier. The precipitation variables have positive coefficients; however, they lose significance in some specifications. The magnitudes of the marginal effects are smaller than those estimated for the hub city pairs. The coefficient on FREEZING is positive but insignificant when we do not include a separate control for the precipitation during cold months, but again becomes negative and significant when we include that control. The major's number of flights at both endpoints now have positive and significant effects on the likelihood of using an owned regional. The marginal effects imply that, at the sample means, an additional flight of the major at the larger endpoint would increase the likelihood of using an owned regional by 0.36 percentage points, while an additional flight at the smaller endpoint would increase the likelihood of using an owned regional by 0.68 percentage points.²⁹ This suggests that on non-hub routes, these variables provide good measures of the integration of the route into the major's network while on hub routes they have no measurable effect. The distance of the route is again insignificant in all of these specifications.

ii. Robustness Checks

We perform several robustness checks on our results. The first set of robustness checks we perform uses alternative measures of weather. In all previous specifications, we have used the maximum of each of our three weather measures across the two endpoints of a city pair. In Table 10, we re-estimate our baseline model (column (2) from Table 6) using three alternate ways of measuring expected weather conditions on a city pair. First, for each city pair, we identify which endpoint airport is the "rainier" endpoint, and calculate all three weather measures for that endpoint. Second, we identify which endpoint airport is the "colder" endpoint and calculate all three weather measures for that endpoint. Finally, we calculate the average of each weather measure across the two end-

²⁹ LARGER_FLIGHTS and SMALLER_FLIGHTS are defined in hundreds of flights.

point airports of a city pair. The results in Table 10 indicate that all three of these perturbations give the same qualitative results.³⁰

Next, we address the issue that our weather variables may be correlated with the unionization of a state in which a regional carrier operates. Specifically, one might worry that our measure of precipitation during cold months could act as a proxy for the greater degree of unionization in the Northeastern states.³¹ We hypothesize that in highly unionized states, the cost advantage of independent regionals (with strong unions) may be smaller than in other states where unions are weaker. Therefore, in the trade-off between control rights and operating costs, the benefit of owning control rights is now more likely to outweigh lower operating costs of independent regionals and therefore make ownership more likely. To investigate this possibility, we determine the state in which each regional airline in our sample has its headquarters. We then include a dummy variable in our regressions that is equal to one if the regional's headquarters are in a more unionized (non-“right to work”) state and equal to zero if the regional is headquartered in a less unionized (“right to work”) state. Even though “right to work” laws do not apply to employees in the airline industry, their existence is likely to reflect the climate of labor relations in the state.^{32 33}

When we include the union dummy we find that it is statistically insignificant in all of our specifications.³⁴ In addition, the inclusion of this variable does not affect the significance or the size of the coefficients on our weather variables. We are therefore confident that the weather effects are not simply proxies for the unionization of the regional airline employees.

Finally, we investigate whether the number of potential competitors on a city pair influences the type of regional that is used on that city pair. We do this as a test of whether foreclosure could be a motive for vertical integration in this industry. We believe, however, that incentives for foreclosure by a major are very small because barriers

³⁰ The same is true when we use these alternate weather measures in the specification that includes only the majors with both types of regionals and dummies for each of these majors.

³¹ We thank George Baker for suggesting this point.

³² “Right to work” laws do not apply to employees in the airline industry because labor relations in this industry are governed by the Railway Labor Act.

³³ Farber (1984) finds that the existence of “right to work” laws is related to preferences against union representation.

³⁴ The results from this exercise are not reported but are available upon request.

to entry are low, especially for existing airlines expanding into new markets. We estimate a two-staged least squares regression in which we instrument for the number of competitors on a city pair with the endpoint population, the number of carriers serving each endpoint³⁵, and a dummy for whether an endpoint is a competitor's hub. The estimation also includes controls for our network variables, the weather variables, and distance. We find that the number of predicted competitors on a city pair has no significant effect on the type of regional used and the coefficients on the network and weather variables have the same signs as in Table 6. This suggests that alternative motivations for vertical integration, such as foreclosure, are unlikely to explain the choice between owned and independent regional carriers.

iii. Airport Level Results

The results so far 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, it is possible that majors do not choose the type of regional to use on a city pair-by-city pair basis but instead decide which type of regional to use at a particular airport and then use that same regional for all city pairs out of that airport which they want to serve with a regional. This would be more 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.

For several reasons, we believe that majors do decide which type of regional to use on a city pair-by-city pair basis. First, we explicitly posed this question to industry participants and they consistently told us that, in their view, there is no good reason not to use multiple regionals at the same airport and that, in fact, this may be preferred because it reduces their dependence on any single regional. Second, empirically, in our sample, we observe 70 instances of the *same* major using both types of regional at a given airport (four of these airports are hubs). Finally, looking at majors who do have multiple regional partners, we often see several of those regionals operating within the same geo-

³⁵ Berry (1992) shows that these variables predict the equilibrium number of carriers serving a route.

graphic regions, suggesting that these regionals could easily serve city pairs out of the same airport.

However, as a further robustness check of our city pair level results, we aggregate the data to the airline-airport level and estimate logit models, with the dependent variable now defined as being equal to one if the airline uses an owned regional at the airport. For this analysis, we restrict our sample to airports at which a major uses only one type of regional. We construct the following airport-level versions of our network variables: AIRPORT_HUB equals one if the airport is a hub to the major, ARRIVAL_HUB equals the fraction of the city pairs served by the carrier from the airport that arrive at one of its hubs, and FLIGHTS equals the number of flights that the major carrier operates out of that airport. PRECIP, FREEZ_PRECIP and FREEZING measure the three weather characteristics for the particular airport.

The results of the airport-level estimation appear in Table 11. The results in columns (1) and (2) of Table 10 indicate that while hub airports are no more likely to be served by owned regionals than non-hubs, airports at which a large fraction of the major's routes go to hubs are more likely to be served by owned regionals. The estimate in column (2) implies that at the sample means a 1 percentage point increase in the number of routes that go from the airport to a hub will increase the likelihood of using an owned regional at the airport by 0.22 percentage points. An increase in the percentage of such routes by one standard deviation will increase the likelihood of using an owned regional by 6.93 percentage points. The finding that the airport-level hub variable is not significant is not surprising given that hubs are such a small fraction of the total number of airports served (though a large fraction of all city pairs have a hub on an endpoint) and that many spoke airports are served by owned regionals.

Columns (3) and (4) add measures of weather at the airport. The results on all of the weather variables are consistent with those obtained earlier. The FREEZING variable remains negative, while both precipitation variables are positive. The hub variable has a positive and significant effect on the likelihood of using an owned regional at the airport when we control for the weather variables in the regression. Finally, we add the carrier's number of flights at the airport in column (5) as another measure of the airport's integration into the major's network. We find no significant effect for this variable.

VI. Conclusion

This paper has developed a simple framework that analyzes the benefits and costs of vertical integration in the U.S regional airline industry and has tested the predictions of this framework using data on major carriers' use of owned and independent regionals in the second quarter of 2000. We have argued that the primary benefit of ownership of a regional is that it provides a major with residual rights of control over how the regional's physical assets and labor force are used. These control rights are valuable because, when unforeseen disruptions occur (as they frequently do in this industry), the major and its regional will generally not agree on how these disruptions should be resolved. In particular, while the major will internalize the impact of any disruption on its overall network, the regional will not. However, offsetting this benefit of vertical integration is the fact that a major's ownership of its regional may partly erode the labor cost savings that are the very reason why majors subcontract to regionals. A major's optimal choice of organizational form will therefore reflect this tradeoff between greater control and lower labor costs.

We have tested this framework by developing two empirical propositions that relate the likelihood of using an owned regional to airline-city pair characteristics that proxy for either the likelihood of unforeseen disruptions or the costs to the major of having these disruptions resolved by the regional. We find empirical support for both of our propositions. Specifically, our results show that owned regionals are more likely to be used on city pairs that are more integrated into the major's network and on city pairs where adverse weather makes schedule disruptions much more likely. We find this to be true even when we look at within-firm variation in the use of owned regionals by controlling for fixed effects for each major airline. These results are robust to alternate ways of constructing our proxies, to various sets of control variables and to analysis at the airline-airport rather than airline-city pair level of observation.

Our paper extends existing empirical tests of the "theory of the firm" in several important ways. First, in contrast to the previous literature, much of which focuses on asset specificity as a determinant of vertical integration, we examine the effect of the costliness and frequency of adaptation decisions on the likelihood of vertical integration.

Second, our paper illustrates the close link between the choice of organizational form and the operational and institutional characteristics of an industry. In our setting, the incentive problems between non-integrated firms arise to a large part out of the network structure of the industry. Similar incentive misalignments may arise in other network industries when firms subcontract a portion of their network to an outside firm. Finally, we present a setting in which part of the costs of vertical integration result from the nature of labor relations in the industry. This suggests the importance of considering not only the costs of vertical integration that derive from lower-powered incentives inside a firm, but also other costs which may result from the unique institutional characteristics and history of an industry.

Appendix A: The link between weather and flight delays

To assess if our measures of the expected weather at an airport indeed capture the frequency of adaptation decisions that need to be made on the city pair or at the airport, we investigate the relationship between *actual* weather and *actual* flight delays. While we do not have direct information on actual flight delays of regional carriers, major airlines are required to report flight delays and early arrivals on a flight-by-flight basis. We use these data for the years 1998-2000 to establish the relationship between flight delays and the weather variables we use in the estimation above. Since major airlines always operate their own flights, 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.³⁶

We have data on the precipitation and the average minimum temperature for each airport in our sample for all months in 1998-2000. For each of these months, we compute average arrival delays for all airline-city pairs between these airports 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 column (1) of Table A.1, we regress average flight delays on the weather variables used in our estimation in Section V. In column (2), we also include airline fixed effects to control for airline-specific factors that influence flight delays. All specifications control for one of the endpoints being the carrier's hub.³⁷

The results in Table A.1 indicate that the relationship we estimate between weather and the likelihood of using an owned regional is very similar to the relationship we estimate between weather and actual delays. Specifically, the estimates suggest that during months with more precipitation, flight delays are longer, especially when the average minimum temperature for the month is below freezing.

While this result may initially be surprising, one way it can be resolved is by considering the geographic pattern of delays that our estimates imply. In particular, our find-

³⁶ A major's decision whether to serve a city pair itself or with a regional should not be correlated with unobservable factors that affect flight delays since this decision is largely based on the size of aircraft that the major wants to use on the city pair.

³⁷ Mayer and Sinai (2003) find that flight delays are more common at hubs.

ings imply that – controlling for other factors, such as hub presence – the longest flight delays are in cold regions with high rain and snowfall, such as the Northeast. The second longest flight delays are in regions with high precipitation but warmer weather, such as the South. These regions tend to have a lot of thunderstorms during the warm months which contribute to disruptions at the airports. The surprising result predicted by our coefficients is that the warm and dry regions of the Southwest have longer flight delays than similarly dry but colder regions such as the Mountain region (this includes states such as Colorado, Idaho, Montana, and Wyoming). We suspect that the explanation for this result is that the population in the Southwest has grown very rapidly during the past decades, whereas the Mountain region has had less rapid population growth. Airport capacity, however, has not grown very much at all in either of these regions. As a result, airports in the Southwest are more likely to be congested than airports in the Mountain states. The negative coefficient estimated on the FREEZING variable (in both the delay and the organizational form regressions) may therefore be capturing this lower level of airport congestion in cold and dry areas.

References:

- Anderson, Erin and David Schmittlein (1984), "Integration of the Sales Force: An Empirical Examination", *Rand Journal of Economics* 15, 385-95.
- Baker, George P. and Thomas N. Hubbard (2004), "Contractibility and Asset Ownership: On-Board Computers and Governance in U.S. Trucking", *Quarterly Journal of Economics* 119(4), 1443-79.
- Ball, Michael, Cynthia Barnhart, George Nemhauser and Amedeo Odoni (2005), "Managing Air Traffic and Airline Operations for Schedule Reliability", *mimeo*, Massachusetts Institute of Technology.
- Berry, Steven T. (1992), "Estimation of a Model of Entry in the Airline Industry", *Econometrica* 60(4), 889-917.
- Borenstein, Severin and Nancy L. Rose (2005), "Regulatory Reform in the Airline Industry", *mimeo*, University of California, Berkeley and Massachusetts Institute of Technology.
- Coase, Ronald (1937), "The Nature of the Firm", *Economica* 4, 386-405.
- Farber, Henry S. (1984), "Right-to-Work Laws and the Extent of Unionization", *Journal of Labor Economics* 2(3), 319-52.
- Gibbons, Robert (2004), "Four Formal(izable) Theories of the Firm?", *mimeo*, Massachusetts Institute of Technology.
- Grossman, Sanford and Oliver Hart (1986), "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Ownership", *Journal of Political Economy* 94: 691-719.
- Hart, Oliver and John Moore (1990), "Property Rights and the Nature of the Firm", *Journal of Political Economy* 98, 1119-58.
- Hubbard, Thomas (2001), "Contractual form and market thickness in trucking", *Rand Journal of Economics* 32(2), 369-86.
- Joskow, Paul (1985), "Vertical Integration and Long-Term Contracts: The Case of Coal-Burning Electric Generation Plants", *Journal of Law, Economics, and Organization* 1, 33-80.
- Joskow, Paul (1987), "Contract Duration and Relationship-Specific Investment: Empirical Evidence from Coal Markets", *American Economic Review* 77, 168-85.
- Lafontaine, Francine (1992), "Agency Theory and Franchising: Some Empirical Results", *Rand Journal of Economics* 23, 263-83.

Lafontaine, Francine and Margaret E. Slade (1997), “Retail Contracting: Theory and Practice”, *Journal of Industrial Economics* 45, 1-25.

Masten, Scott E. (1984), “The Organization of Production: Evidence from the Aerospace Industry”, *Journal of Law and Economics* 27, 403–17.

Masten, Scott E. and Keith Crocker (1985), “Efficient Adaptation in Long-Term Contracts: Take-or-Pay Provisions for Natural Gas”, *American Economic Review* 75, 1083–93.

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.

Monteverde, Kirk, and David Teece (1982), “Supplier Switching Costs and Vertical Integration in the Automobile Industry”, *Bell Journal of Economics* 13, 206-12.

Nickerson, Jack 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, 433-37.

Williamson, Oliver (1971), “The Vertical Integration of Production: Market Failure Considerations”, *American Economic Review* 63, 112-23.

Williamson, Oliver (1975), *Markets and Hierarchies: Analysis and Antitrust Implications*, New York, NY: Free Press.

Williamson, Oliver (1985), *The Economic Institutions of Capitalism*, New York, NY: Free Press.

Table 1
Regional Airline Statistics

	2000
Carriers Operating	94
Passengers Enplaned (millions)	84.6
Revenue Passenger Miles (billions)	25.27
Average RPM's per Carrier (millions)	268.8
Available Seat Miles (billions)	42.55
Average Load Factor (percent)	59.39
Departures Completed (millions)	4.46
Airports Served (North America)	729
Aircraft Operated	2,271
Average Passenger Trip Length (miles)	299
Average Seating Capacity (seats/aircraft)	31.7

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

Table 2
Majors and Regional Partners in 2000
Regional carriers in bold are fully owned by the major

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/CCAair Chautauqua Airlines Colgan Airways Commutair Mesa Air Group/Mesa Airlines Piedmont Airlines PSA Airlines

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

Table 3
Variable Names and Definitions

VARIABLE NAME	DEFINITION	SOURCE
OWNED_REGIONAL	=0 if airline operates flight through an independent regional partner, =1 if airline operates flight through a regional partner which it owns	OAG and RAA
HUB	=1 if either endpoint is carrier's hub	Authors' construction
FREEZING	= Average # of months per year in which average daily minimum temperature is below 0 Celsius; maximum of the two endpoint airports of a city pair (based on 1970-1995 data)	Spatial Climate Analysis Service
PRECIP	= Average annual precipitation, in millimeters; maximum of the two endpoint airports of a city pair (based on 1970-1995 data)	Spatial Climate Analysis Service
FREEZ_PRECIP	= Average annual precipitation, in millimeters, during months in which the average daily minimum temperature is below 0 Celsius; maximum of the two endpoint airports of a city pair (based on 1970-1995 data)	Spatial Climate Analysis Service
LARGER_FLIGHTS	= Major's # of departing domestic flights on other city pairs from endpoint at which it is larger, in hundreds	OAG
SMALLER_FLIGHTS	= Major's # of departing domestic flights on other city pairs from endpoint at which it is smaller, in hundreds	OAG
DISTANCE	= Distance of the route (in hundreds of miles)	Authors' construction

Table 4
Summary Statistics
Routes Served by Either Type of Regional

VARIABLE NAME	N	MEAN	ST. DEV.	MIN	MAX
OWNED_REGIONAL	994	0.55	0.50	0	1
HUB	994	0.70	0.46	0	1
LARGER_FLIGHTS (00s)	994	2.04	1.56	0	6.18
SMALLER_FLIGHTS (00s)	994	0.06	0.16	0	3.23
FREEZING	994	3.28	1.72	0	8.54
PRECIP	994	1104.95	273.34	123.16	1738.088
FREEZ_PRECIP	994	207.49	112.88	0	462.96
DISTANCE (00s of miles)	994	3.14	2.06	0.20	10.76

Table 5A
Characteristics of Routes (<1500 miles) Served with Regionals (any type)
Means of Selected Variables

VARIABLE	MAJOR (N=1017)	REGIONAL (N=994)
HUB	0.81	0.70
LARGER_FLIGHTS (00s)	3.88	2.04
SMALLER_FLIGHTS (00s)	0.303	0.061
FREEZING	3.26	3.28
PRECIP	1130.16	1104.95
FREEZ_PRECIP	196.07	207.49
MEAN_POP	4,392,790	3,991,375
DISTANCE (00s of miles)	7.30	3.14

Table 5B
Characteristics of Routes Served by Owned vs. Independent Regionals
Means of Selected Variables

VARIABLE	OWNED (N=547)	INDEPENDENT (N=447)
HUB	0.78	0.61
LARGER_FLIGHTS (00s)	2.26	1.76
SMALLER_FLIGHTS (00s)	0.066	0.054
FREEZING	2.93	3.72
PRECIP	1188.17	1003.12
FREEZ_PRECIP	206.69	208.48
MEAN_POP	3,944,620	4,062,449
DISTANCE (00s of miles)	3.32	2.93

Table 6
Major's Choice Between Owned and Independent Regional: Logit (All Carriers)
Marginal Effects

	Dependent Variable = OWNED_REGIONAL				
	(1)	(2)	(3)	(4)	(5)
HUB	0.1990 (0.0349)**	0.3272 (0.0376)**	0.3892 (0.0371)**	0.3750 (0.0502)**	
FREEZING		-0.0409 (0.0119)**	-0.1582 (0.0208)**	-0.1586 (0.0216)**	-0.1461 (0.0205)**
PRECIP		0.0008 (0.0001)**	0.0005 (0.0001)**	0.0005 (0.0001)**	0.0004 (0.0001)**
FREEZ_PRECIP			0.0019 (0.0003)**	0.0020 (0.0003)**	0.0018 (0.0003)**
LARGER_FLIGHTS				0.0138 (0.0154)	0.0897 (0.0129)**
SMALLER_FLIGHTS				0.2122 (0.2663)	0.1087 (0.1333)
DISTANCE	0.0144 (0.0082)+	-0.0013 (0.0089)	0.0067 (0.0093)	0.0058 (0.0094)	0.0136 (0.0091)
Carrier fixed effects	No	No	No	No	No
Observations	994	994	994	994	994

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 7
Major's Choice Between Owned and Independent Regional: Logit
(Carriers with Both Types, Including Carrier Dummies)
Marginal Effects

	Dependent Variable = OWNED_REGIONAL				
	(1)	(2)	(3)	(4)	(5)
HUB	0.4321 (0.0492)**	0.5865 (0.0511)**	0.6230 (0.0518)**	0.6843 (0.0549)**	
FREEZING		0.0596 (0.0210)**	-0.0986 (0.0350)**	-0.0981 (0.0361)**	-0.0462 (0.0341)
PRECIP		0.0013 (0.0001)**	0.0009 (0.0002)**	0.0010 (0.0002)**	0.0008 (0.0002)**
FREEZ_PRECIP			0.0020 (0.0004)**	0.0021 (0.0004)**	0.0013 (0.0004)**
LARGER_FLIGHTS				-0.0316 (0.0198)	0.1620 (0.0234)**
SMALLER_FLIGHTS				0.2926 (0.2117)	0.1343 (0.1022)
DISTANCE	-0.0035 (0.0104)	-0.0257 (0.0120)*	-0.0206 (0.0128)	-0.0194 (0.0132)	-0.0195 (0.0110)+
Carrier fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	699	699	699	699	699

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 8
Major's Choice Between Owned and Independent Regional: Logit (All Carriers)
Hub Routes
Marginal Effects

	Dependent Variable = OWNED_REGIONAL		
	(1)	(2)	(3)
FREEZING	-0.0969 (0.0141)**	-0.1317 (0.0221)**	-0.1236 (0.0223)**
PRECIP	0.0009 (0.0001)**	0.0008 (0.0001)**	0.0009 (0.0001)**
FREEZ_PRECIP		0.0007 (0.0004)*	0.0005 (0.0004)
LARGER_FLIGHTS			-0.0246 (0.0160)
SMALLER_FLIGHTS			0.0203 (0.0867)
DISTANCE	-0.0058 (0.0102)	-0.0021 (0.0104)	-0.0019 (0.0104)
Carrier fixed effects	No	No	No
Observations	700	700	700

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 9
Major's Choice Between Owned and Independent Regional: Logit (All Carriers)
Non-Hub Routes
Marginal Effects

	Dependent Variable = OWNED_REGIONAL		
	(1)	(2)	(3)
FREEZING	0.0305 (0.0223)	-0.1368 (0.0437)**	-0.0862 (0.0462)+
PRECIP	0.0004 (0.0002)*	0.0002 (0.0002)	0.0003 (0.0002)
FREEZ_PRECIP		0.0021 (0.0005)**	0.0014 (0.0005)**
LARGER_FLIGHTS			0.3607 (0.0994)**
SMALLER_FLIGHTS			0.6754 (0.3244)*
DISTANCE	0.0157 (0.0173)	0.0132 (0.0182)	-0.0151 (0.0238)
Carrier fixed effects	No	No	No
Observations	294	294	294

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 10
Major's Choice Between Owned and Independent Regional: Logit (All Carriers)
Alternate Weather Measures
Marginal Effects

	Dependent Variable = OWNED_REGIONAL		
	(1)	(2)	(3)
HUB	0.4003 (0.0367)**	0.3857 (0.0366)**	0.3842 (0.0374)**
<i>Mean of Weather Variables Across Endpoints</i>			
FREEZING	-0.2174 (0.0270)**		
PRECIP	0.0005 (0.0001)**		
FREEZ_PRECIP	0.0026 (0.0004)**		
<i>Weather Measured at Rainer Endpoint</i>			
FREEZING		-0.2118 (0.0277)**	
PRECIP		0.0005 (0.0001)**	
FREEZ_PRECIP		0.0026 (0.0004)**	
<i>Weather Measured at Colder Endpoint</i>			
FREEZING			-0.1462 (0.0214)**
PRECIP			0.0004 (0.0001)**
FREEZ_PRECIP			0.0017 (0.0003)**
DISTANCE	0.0100 (0.0093)	-0.0087 (0.0097)	0.0244 (0.0090)**
Carrier fixed effects	No	No	No
Observations	994	994	994

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Table 11
Major's Choice Between Owned and Independent Regional: Logit (All Carriers)
Airport Level Regressions - All Airports
Marginal Effects

	Dependent Variable = OWNED_REGIONAL				
	(1)	(2)	(3)	(4)	(5)
AIRPORT_HUB	-0.0173 (0.1164)	0.1460 (0.1119)	0.2405 (0.0890)**	0.2509 (0.0867)**	0.2690 (0.1402)+
ARRIVAL_HUB		0.2189 (0.0699)**	0.3373 (0.0849)**	0.3581 (0.0819)**	0.3566 (0.0824)**
FREEZING			-0.0523 (0.0118)**	-0.0956 (0.0212)**	-0.0956 (0.0212)**
PRECIP			0.0004 (0.0001)**	0.0003 (0.0001)**	0.0004 (0.0001)**
FREEZ_PRECIP				0.0008 (0.0003)*	0.0008 (0.0003)*
FLIGHTS					-0.0114 (0.0686)
Carrier fixed effects	No	No	No	No	No
Observations	626	626	626	626	626

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Table A.1
The Effect of Weather on Flight Delays: OLS (Major Carriers)
City Pairs between Top 300 Airports Served by Major Carriers

	Dependent Variable = Average Arrival Delay	
	(1)	(2)
FREEZING	-1.4079 (0.1221)**	-1.2014 (0.1219)**
PRECIP	0.0101 (0.0004)**	0.0140 (0.0004)**
FREEZ_PRECIP	0.0302 (0.0017)**	0.0310 (0.0017)**
HUB	1.2709 (0.0679)**	0.6647 (0.0829)**
Carrier fixed effects	No	Yes
Observations	130,284	130,284

Robust standard errors in parentheses
+ significant at 10%; * significant at 5%; ** significant at 1%