Does Vertical Integration Affect Firm Performance? Evidence from the Airline Industry

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: August 2006
This Draft: July 2007

Abstract

This paper investigates the performance implications of firms' vertical integration decisions. Our setting is the U.S. airline industry. Major airlines subcontract service on low-density short and medium-haul routes to regional airlines. These regional partners are either owned by the major airline or are independently owned and contract with one or more major airlines. Earlier work (Forbes and Lederman, 2007) has argued that the primary benefit of ownership of a regional is that it mitigates incentive problems that arise when unforeseen schedule disruptions require the major to make changes that involve its regional partner's operations. We explicitly test this hypothesis by estimating the relationship between a major's performance on flights that it operates out of a given airport and its extent of vertical integration with the regional partners that operate flights on its behalf out of that airport. Our estimation approach accounts for the endogeneity of airlines' ownership decisions. In addition to estimating an overall performance effect, we also investigate how this effect changes as the likelihood of schedule disruptions increases. We exploit the fact that while ownership decisions are fixed in the short-run, an important source of unanticipated schedule disruptions – weather - changes on a day-to-day basis. Because ownership decisions cannot respond to these daily changes, we are able to observe carriers with both owned and independent regionals in a variety of weather conditions. Our results indicate that majors using owned regionals at an airport experience shorter delays and fewer cancellations than majors at the same airport using independent regionals. Moreover, this performance advantage increases as weather deteriorates. The effects we estimate are both statistically and economically significant. Our results provide one of the first pieces of direct evidence that contracts cannot replicate the incentive alignment achieved through ownership.

I. Introduction

Economists have long been interested in the question of what determines firm boundaries. There is now a large theoretical literature that seeks to explain which activities are optimally carried out inside firms and which are more efficiently carried out through market transactions.[1] Much of this literature emphasizes the role of incomplete contracts. If complete contracts cannot be written, then there is a role for ownership to mitigate various incentive problems that can arise when transactions are governed by incomplete contracts. More recently, a growing empirical literature that seeks to test this "theory of firm" has developed. Researchers typically identify transaction characteristics that proxy for the magnitude of the incentive problems that arise when contracts are incomplete and test whether transactions with these characteristics are more likely to be carried out in-house.[2] An affirmative finding is interpreted as evidence that vertical integration does indeed mitigate incentive problems that arise under incomplete contracts.

Note, however, that this empirical approach provides only an *indirect* test of the costs and benefits of vertical integration hypothesized in the theory. Moreover, it cannot shed light on the magnitude of these costs and benefits. An alternate approach is to look for evidence that the performance of integrated and non-integrated firms differs on specifically those margins hypothesized in the theory. In addition to providing *direct* evidence on the costs and benefits of vertical integration, this approach also allows one to quantify them. For example, if the theory predicts that a benefit of vertical integration is to improve the coordination between upstream and downstream production units, while a cost of vertical integration is that employees receive lower-powered incentives, an indirect test of this theory would be to identify transactions for which the benefits of

---

[1] At the broadest level, three theoretical perspectives can be considered to co-exist in this literature: agency theory (which emphasizes the tradeoff between insurance and optimal incentive provision; e.g. Alchian and Demsetz, 1972, and Holmström, 1982), transaction costs theory (which emphasizes the role of firms, or hierarchies, in mitigating hold-up problems; see Williamson, 1975, 1985) and property rights theory (which emphasizes the role of asset ownership in providing residual rights of control which, in turn, influence ex ante investment decisions; see Grossman and Hart, 1986, and Hart and Moore, 1990).
[2] This approach has frequently been used to test the effect of asset specificity or complexity on the likelihood that a transaction is organized internally; see, for example, Monteverde and Teece (1982), Anderson and Schmittlein (1984), Masten (1984), Masten and Crocker (1985), Joskow (1985, 1987), and Hubbard (2001). Although less frequently, a similar approach has been used to test predictions of the agency model as well the property-rights model (see Lafontaine and Slade, 2007, for a review of the literature).

coordination are presumed to be particularly large or the costs of low-powered incentives particularly low and then test whether these transactions are more likely to be vertically integrated. A direct test of the theory would instead try to measure whether transactions that are carried out in-house are better coordinated or suffer from suboptimal effort of employees.

Despite their advantages, studies providing direct evidence on the performance consequences of vertical integration decisions are rare. Such evidence has been difficult to establish empirically, both because obtaining data on the relevant performance measures is difficult and because one must account for the fact that vertical integration decisions are endogenous (Masten, 1993). This paper overcomes these difficulties and provides one of the first pieces of direct evidence that firms' vertical integration decisions do, indeed, affect specific measures of their performance. We estimate the effect of vertical integration on operational performance in the U.S. airline industry. The large U.S. network carriers, often called "majors", employ regional airlines to operate flights on low-density short and medium-haul routes. These flights are operated under codeshare agreements such that the regional operates flights on behalf of the major carrier, who markets and tickets these flights under its own flight designator code. 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. There is substantial variation, both across and within airlines, in the use of owned and independent regional airlines.

Our analysis builds on Forbes and Lederman (2007a, hereafter FL) which analyzes the benefits and costs of vertical integration between major and regional airlines. As they explain, incentive problems between majors and independently owned regional partners arise when schedule disruptions occur and real-time adjustments to the schedule set by the major and the regional in their contract are required. While the major airline's optimal response to a disruption will internalize the impact of the disruption on its entire network, an independent regional instead has an incentive to optimize its response only over the routes that it serves for the major. This may cause an independent regional to take actions that are not in the best interest of the major. Ownership of a

regional mitigates this incentive problem by giving the major residual rights of control over how the regional's assets are used. FL explain that majors will use owned regionals when these benefits of vertical integration outweigh the higher labor costs that may result.

We test the framework developed in FL by relating a major's performance on flights that *it operates itself* to its extent of vertical integration with regional partners who operate *other* flights out of the same airport. Note that this setup – which relates the performance on one set of transactions to a firm's vertical integration decision on a different set of transactions – follows naturally from the network structure of this industry. In particular, the fact that ownership allows a major to internalize the impact of schedule disruptions on its entire network implies that vertical integration between a major and its regional will have an impact not only on the performance on flights operated by the regional, but also on the performance on flights operated by the major itself. It is this second effect that we focus on.

Our empirical setting provides us with three important benefits. First, we have access to two detailed performance measures – flight delays and cancellations – which are closely related to the overall performance of an airline's network. Second, our data allow us not only to estimate the average relationship between ownership and performance but they also let us trace out how this relationship changes with the transaction environment. We observe the same transaction (in our case, flight) every day over the course of a whole year. While an airline's vertical integration decision is fixed in the short-term, the likelihood of schedule disruptions changes from day to day as the weather conditions at an airport change. As a result, we observe carriers with both owned and independent regionals on days with "good" weather conditions – when schedule disruptions are less likely – and on days with "bad" weather conditions – when schedule disruptions are particularly likely. We believe that we are the first to estimate how the performance effects of vertical integration vary with the transaction environment by observing both integrated and non-integrated transactions in a range of transaction environments. This approach provides a novel test of the theory and a new source of identification.

Finally, we are able to instrument for the endogenous vertical integration decision. Our instruments are derived from the findings in FL which shows that majors are more likely to use owned regionals on routes that are more integrated into the major airline's network, and on routes that experience more adverse weather on average. We instrument for a major's extent of vertical integration with its regional partners at a given airport with the characteristics of *other* endpoint airports served by the regional partners.

Our main specification estimates the relationship between a major's departure delay on a flight and the fraction of its regional flights at the departure airport of that flight which are operated by an owned regional partner. In addition to estimating the direct effect of the extent of vertical integration at the airport on the airline's performance, we also interact our vertical integration variable with measures of adverse weather. In particular, we have daily data on precipitation and temperatures at all airports in our sample. We divide our sample into a "summer sample" and a "winter sample" so that we can separately estimate the effects of rain and snow in the months in which these two types of adverse weather most commonly occur.

Our results strongly indicate that vertical integration does, indeed, improve airlines' network performance and, even more so, in the presence of adverse weather. For example, the results from our summer sample suggest that, on days with "good" weather, majors using only owned regionals at an airport experience flight delays that are approximately five minutes shorter than those experienced by majors using only independent regionals. On days with "bad" weather, majors using only owned regionals have a performance advantage that translates into delays that are shorter by 13 minutes on average. The results using cancellations as the dependent variable and the results from the winter sample are qualitatively similar. These findings provide direct evidence that, at least in our setting, contracts are not able to replicate the incentive alignment that is achieved by a major's ownership of its regional.

Our paper is most closely related to two recent studies that also test predictions from the "theory of the firm" literature by estimating the effect of organizational form on performance. Novak and Stern (forthcoming) investigate how automobile manufacturers' vertical integration decisions affect two specific performance margins in automobile product development. Consistent with predictions from the theoretical literature, they

find that integration is associated with lower initial performance but higher performance improvement over the life of the automobile model. Ciliberto (2006) investigates whether hospitals that have vertically integrated with their physicians invest more than hospitals than negotiate managed care contracts independently of their doctors. If contracts between hospitals and doctors are complete, then vertical integration should have no impact on investment incentives. He finds that hospitals that are vertically integrated do add more healthcare services over time suggesting that contracts cannot fully replicate the incentives of integration.

Our paper is also related to a small literature that looks at the overall performance consequences of choosing an organizational form that is inconsistent with the transaction environment. This literature on so-called "transactional misalignment" originates with Masten *et al.* (1991) and includes a number of recent contributions primarily from the strategy field. Finally, this paper is also related to a small literature that investigates the effects of vertical integration on a host of other outcome variables such as prices, costs and quantities (see Lafontaine and Slade, 2007, for a summary).

We believe that this paper contributes to the literature in two important ways. First, it is the first to trace out how the relationship between vertical integration and performance changes with the transaction environment by exploiting high frequency data on performance and on variables that affect the returns to vertical integration. A similar approach may be fruitful in other settings in which adjustment costs or institutional factors prevent organizational form decisions from changing as quickly as transaction characteristics change.[3] The benefit of this approach is that it provides rich identification of the variables of interest from the data, as opposed to relying on functional form assumptions.

Second, our results show that a firm's vertical integration decision on one set of transactions can have implications for its performance on a different set of transactions. In our case, these effects result from the fact that there are externalities across transactions that are organized in a network. We believe that similar effects may be

---

[3] Ahmadjian and Oxley (2007) use a related approach to estimate the relationship between "hostage" arrangements and supplier performance. They exploit the fact that equity ties between Japanese buyers and suppliers are generally fixed over time and then investigate how the presence of an equity tie affects supplier performance as demand conditions change.

present in other industries when firms try to outsource a fraction of their network and, more generally, when there are externalities across transactions.

The remainder of the paper is organized as follows. Section II presents the industry background and theoretical considerations. Section III describes our empirical approach. Section IV addresses data and measurement issues. We present our results in Section V and offer concluding thoughts in Section VI.

II. Theoretical Considerations

II.A. Background: The Role of Regional Airlines[4]

Regional airlines operate as "subcontractors" for major U.S. network carriers on low-density short and medium-haul routes. These routes are most efficiently served with the small regional jets and turbo-prop planes that regional airlines operate. Majors subcontract this service to regional airlines because regionals have a cost advantage on these types of routes. This cost advantage results primarily from the lower compensation that regional airline employees receive.[5] Majors do not operate any flights of their own with regional aircraft.

Regional airlines operate under codeshare agreements with one or more major carriers. 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, Delta's regional Comair operates under the name Delta Connection). To facilitate passenger connections between a major and its regional, their schedules, as well as check-in and baggage handling, are coordinated.

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.[6] In the case of an owned regional, "vertical integration"

---

[4] For a detailed description of the role of regionals in the U.S. airline industry, we refer the reader to Forbes and Lederman (2007b).

[5] See Forbes and Lederman (2007a) for a detailed discussion of the source of lower labor costs among regional airline employees.

[6] Theoretically, it would be possible for owned regionals to perform contract flying for other majors, but we do not observe that in the data. We believe that this is due to the risk of a hold-up problem that would arise if an owned regional performed services for a competitor of the major who owns it.

means that the major carrier owns the assets of the regional but the regional and the major maintain separate operations and labor contracts.[7]

Table 1 lists the major-regional partnerships that were in place in 2000 (the year of our sample) 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.

II.B. The Benefits and Costs of Vertical Integration

The relationship between vertical integration and performance that we examine follows from the framework developed in FL, which investigates the determinants of vertical integration between a major and a regional. As FL explain, incentive problems between major airlines and their independently owned regional partners arise when real-time adjustments to the regional's planned schedule are necessary. The need for such adjustments can arise for a large number of reasons, such as mechanical problems, adverse weather, or security disruptions. While the major airline's optimal response to a disruption will attempt to internalize the impact of the disruption on its entire network, an independent regional instead has an incentive to optimize its response only over the routes that it serves for the major. For example, when adverse weather forces an airline to reduce the number of its takeoffs and landings, a major may prefer to cancel several of its (low-capacity) regional flights to allow more of its (high-capacity) mainline flights to operate. However, the regional – which is compensated only based on the routes it serves – will not.

---

[7] 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.

Note that this incentive problem occurs because there are externalities across the major's and the regional's flights that arise for two reasons: (1) the integration of their flights into a common network (which results in flows of passengers and cargo between their planes); and (2) the fact that their flights compete for scarce inputs such as takeoff and landing slots or airport gates. As a result of these externalities, an independent regional will not resolve unforeseen contingencies in a manner that internalizes the impact of its actions on the remainder of the major's network. Ownership of a regional mitigates this incentive problem by giving the major residual rights of control over how the regional's assets are used.[8] This allows the major to respond to unforeseen schedule disruptions in way that re-optimizes its overall network. Note that this implies that vertical integration between a major and its regional partner will have an impact not only on the performance of the flights operated by its regional, but also on the performance of the flights that the major operates itself.

Offsetting these benefits of vertical integration are costs that are associated with a major's ownership of its regional. These costs of ownership stem primarily from the higher labor costs that can arise when a regional is owned. A major will choose to use an owned regional on routes on which the benefits of ownership are large enough to offset these costs. FL test this framework by estimating whether owned regionals are more likely to be used on routes on which the probability of adaptation decisions is higher and on routes on which the costs of having adaptations decisions resolved sub-optimally are higher. They measure the former using average weather patterns and the latter with measures of a route's integration into the major's network. They find empirical support for both hypotheses.

III. Empirical Approach

III.A. General Identification Issues

---

[8] As FL describe, contracts between majors and regionals try to address this incentive problem by allocating to the major the right to make changes to the regional's planned schedule. However, this does not fully solve the incentive problem because the schedule changes ordered by the major are still carried out by the regional. Thus, the regional is left with residual rights of control over issues which cannot be explicitly contracted on (for example, the speed with which the regional's flights are cancelled) and may use these rights to make decisions that are not optimal for the major.

Estimating the effects of organizational form on firm performance presents an empirical challenge because organizational form decisions are endogenously determined by optimizing firms. Specifically, as Coase (1937) and Williamson (1975) argue, firms will choose the organizational form that minimizes transaction costs. Figure 1 (taken from Williamson, 1991) illustrates their basic argument: although transaction costs increase for firms as well as for markets as transaction difficulty, k, increases, they increase at a faster rate for markets than for firms (Williamson, 1991). As a result, there is a critical value of transaction difficulty, k^*, above which firms are the optimal organizational form and below which markets are the optimal organizational form.

Figure 1 also illustrates the two empirical issues that confront researchers trying to estimate the performance effects of organizational form decisions. First, as Masten (1993) and Gibbons (2004) have emphasized, a simple comparison of the performance of firms and markets will be misleading because markets are chosen for transactions with low levels of difficulty while firms are chosen for transactions with high levels of difficulty. Thus, the firms that are observed may appear to be less efficient than the markets that are observed when, in fact, the firms that are observed may be more efficient than the markets they replace. Econometrically, this amounts to a sample selection problem which induces correlation between the organizational form variable and the error term in the performance equation. It can be resolved by carrying out a Heckman selection correction or by using an instrumental variables approach.

Second, as Masten also points out, the effect of organizational form on performance will typically vary with transaction characteristics. This is the case in Figure 1 where both the intercepts and the slopes of two curves are different and therefore the performance effect is different, and even switches signs, at different levels of k. Accurately estimating performance effects therefore requires one to estimate not only the average effect but also the change in this effect as transaction difficulty varies. However, when each organizational form is only observed in those settings in which it is optimal (i.e.: when markets are only observed for $k<k^*$ and firms for $k>k^*$), estimates of the effect of organizational form on performance rely on functional form assumptions for identification.

Figure 1

III.B. The Empirical Setting

We now turn to the details of our empirical approach. As explained in Section II, the theoretical prediction that we seek to test is whether majors using owned regionals are able to better optimize their overall network, in particular when unforeseen schedule adjustments are necessary. Schedules are typically set assuming good flying conditions. However, disruptions can occur for a large number of reasons – for example, due to adverse weather, air traffic control problems, security breaches or mechanical problems. When such disruptions occur, an airline may be forced to delay or cancel some of its own or its regional's flights. It will do so with the objective of minimizing the impact of the schedule disruptions on its profits. While we clearly cannot observe the profit implications of the airline's schedule adjustments, we can observe the actual adjustments that are made. Specifically, we have flight level data on delays and cancellations. Thus, our empirical approach assumes that longer delays and more frequent cancellations have a negative effect on profits. We know from the results in Forbes (2007) that flight delays indeed have a negative effect on route-level fares and profits.[9]

---

[9] In addition, the results in Rupp and Holmes (forthcoming) show that flight cancellations are negatively correlated with route-level revenues.

We estimate the impact of a major's use of an owned regional at a particular airport on its performance on *other* routes out of that same airport that the major serves itself. We have variation both across majors and within some majors in their extent of owned regional use. In contrast, the approach taken in existing studies is to estimate the effects of vertical integration of a given transaction on a firm's performance on that same transaction. We adopt our setup because it is consistent with the theoretical framework summarized in Section II. As explained there, ownership of its regional allows the major to respond to schedule disruptions in a way that maximizes the performance of its overall network. As a result, a major's vertical integration decision will have an impact not only on the performance of the flights operated by its regional, but also on the performance of the major's own flights. It is this second effect that we focus on.

Note that we only observe the performance of the major's own flights, not the performance of its regionals' flights.[10] One might be concerned that majors with owned regionals tradeoff fewer delays and cancellations on their own flights for more delays and cancellations on their regionals' flights – something that may be harder to do for majors with independent regionals. If this were true, then our assumption that lower delays (and fewer cancellations) represent a more optimized network may not be true since we would not be accounting for the delays or cancellations on the major's flights operated by regional partners. However, given the much larger size of the major's operations relative to the regional's, this effect is unlikely to offset the performance effects we find for the major's flights.[11]

A unique feature of our setting is the panel structure of our data. We observe the same transaction (in our case, flight) every day over the course of a whole year. While an airline's vertical integration decision is fixed in the short-term due to the adjustment costs of changing its organizational form, the returns to vertical integration at an airport change from day to day as the weather conditions - and with them the likelihood of unanticipated schedule disruptions - change. As a result, in our setting, we observe carriers with both owned and independent regionals on days with "good" and with "bad" weather

---

[10] Regional airlines have only recently started to report flight delays, as some of them have become large enough to meet DOT reporting requirements.

[11] In the year of our sample period, one in seven U.S. domestic passengers traveled on a regional airline.

conditions. In other words, we observe both types of organizational form over the full range of transaction difficulty.

Our setup also provides us with a set of instrumental variables for a major's vertical integration decision. These instruments are derived from the findings in FL. As described above, that study shows that majors are more likely to use owned regionals on routes that are more integrated into the major airline's network, and on routes that experience more adverse weather on average. Because we are testing whether a major's departure delays and cancellations on a route are affected by the major's vertical integration with a regional partner that serves *other* routes from the *same* departure airport, we use the characteristics of the other airports served by the regional partner from this departure airport as instruments for the major's decision to vertically integrate with the regional. For example, assume that one of the routes in our sample is Boston-Atlanta and the major carrier who serves the route is Delta. We are interested in testing if Delta's decision to own or contract with the regional carrier that it uses to serve other routes out of Boston – such as Boston-Burlington or Boston-Syracuse – affects Delta's performance on the Boston-Atlanta route. We use the average weather conditions and network integration of the Burlington and Syracuse airports as instruments for Delta's extent of vertical integration at the Boston airport.

We require that our instruments be correlated with an airline's vertical integration decision on the routes served by its regionals, but be uncorrelated with the error of the performance equation on the routes served by the major itself. Since delays can propagate through an airline's network, it is important for our use of these instruments that characteristics of other airports served by the regional do not contribute unobserved shocks to delays on the major's own route. Continuing with the example used above, we must assume that average weather conditions at Burlington and Syracuse do not affect Delta's delays on the Boston-Atlanta route. While this might initially seem like an ambitious assumption, note that aircraft and crew are not shared across majors and regionals. Thus, the primary mechanism through which delays propagate throughout a network – aircraft and crew not being where they need to be – simply does not operate here.

13

Another reason why delays may propagate is that planes may wait for connecting passengers who arrive late. Since passengers do connect from regionals to majors, this could create correlation between our instruments and the error in the performance equation. However, note that, in general, airlines will only delay the last flight of the day to wait for late connecting passengers because it is costly for airlines, and very damaging to their customer satisfaction, if passengers have to stay at a connecting airport overnight. To account for this possible source of correlation, we estimate specifications that explicitly control for whether a flight is the last one of the day on a route and we find that our results are robust to this check.

Finally, there may be some remaining correlation between characteristics of other airports served by the regional and the error in the performance equation that is due to shared airport facilities at the departure airport. If such correlation exists, it would lead our results to be biased towards finding that integrated airlines perform worse, not better, since FL show that airlines vertically integrate on routes that are *more* likely to experience schedule disruptions. Our findings that vertically integrated airlines perform better than non-integrated airlines reassure us that possible correlation between our instruments and the error in the performance equation is not a great concern for our results.

III.C Estimation Equation

We estimate the following performance equation:

$$P_{ir}^t = X_{ir}^t \beta + Z_{ir}^t \gamma + \delta_1 VI_{ir} + Z_{ir}^t * VI_{ir} \delta_2 + \varepsilon_{ir}^t \qquad (1)$$

where P_{ir}^t is a measure of the performance of airline i on flight r on day t, X_{ir}^t and Z_{ir}^t are vectors of variables that affect airline i's performance on flight r on day t, VI_{ir} is a scalar that captures airline i's vertical integration decision at the origin airport of flight r, and ε_{ir}^t is an error term. The difference between X_{ir}^t and Z_{ir}^t is that the effect of the former is independent of whether firm i is vertically integrated, whereas the latter may have a differential effect on performance depending on whether airline i is vertically integrated or not. In our setting, Z_{ir}^t contains measures of daily weather (with higher values of Z_{ir}^t capturing worse weather) while X_{ir}^t contains other controls which affect

delays or cancellations. As described above, we control for the endogeneity of the vertical integration decision by instrumenting for VI_{ir} and $Z_{ir}^t * VI_{ir}$.

Our hypothesis tests focus on δ_1 - which measures the difference in performance between majors using owned and independent regionals in "good" weather - and on δ_2 - which measures how this performance difference changes as weather conditions deteriorate. If ownership of a regional allows a major to better optimize its network, then we expect both δ_1 and δ_2 to be less than zero.

IV. Data and Measurement

IV.A. Data Sources

Our empirical analysis is based on several sources of data. The primary source is flight-level on-time statistics from the U.S. Bureau of Transportation Statistics. This database contains every flight operated by all major U.S. carriers.[12] Each observation in the data corresponds to a particular flight on a particular day and contains information on the operating carrier, the departure and arrival airports, the scheduled and actual departure and arrival times, the time spent on the runway at the departure and the arrival airport, and whether the flight was cancelled or diverted.

We augment these data with information from several other sources. First, data from the Official Airline Guide (OAG) provide us with the complete flight schedules of all domestic airlines, regionals as well as majors.[13] The OAG data allow us to measure an airline's total scale of operations as well as the scale of operations of each of its regional partners, at each airport at which it operates. We combine these data with information from the Regional Airline Association (RAA) that shows which regional airlines are owned by a particular major. Together, the OAG and RAA data allow us to calculate an airline's extent of vertical integration with its regionals at each airport at which it operates. Finally, data on the daily weather at each airport are taken from the National Oceanographic and Atmospheric Administration (NOAA).

---

[12] Carriers are required to report these data if they account for at least one percent of domestic passenger revenues in the prior year.

[13] Our data provide a representative week for each quarter.

IV.B. Construction of the Sample

Our sample year is 2000. The sample includes all flights operated by the seven large network carriers (American, Continental, Delta, Northwest, TWA, United and US Airways).[14] We divide our sample into a "summer sample" and a "winter sample" so that we can separately estimate the effects of rain and snow in the months in which these two types of adverse weather most commonly occur. Our summer sample includes all routes but is restricted to the months of May to October. We also exclude any days that fall between May and October on which the average temperature is below freezing. Our winter sample includes only routes that depart from airports in the Northeast or Midwest census regions and is restricted to flights that operate between December and February.

We impose several additional restrictions on both the summer and winter samples. First, we drop observations for which our weather data are missing. Second, we exclude from each sample airports which have flights departing from them for less than 50% of the days in that sample.[15] Third, we exclude airports that are only served by one of the seven airlines in our sample because, once we include departure airport-date fixed effects, these airports do not help identify our vertical integration variables. Fourth, we exclude routes to or from New York's LaGuardia airport because LaGuardia changed its slot control rules during 2000, resulting in a large increase in delays (see Forbes, 2007, for details). Fifth, because we are relating a major's departure delay on a route to its vertical integration with a regional at the departure airport, we exclude routes that depart from an airport at which the major does not use a regional at all. Finally, we exclude flights on Saturdays and Sundays so that the variation in an airline's extent of vertical integration is not simply capturing within-week fluctuations in regional use. After imposing these restrictions, our final summer dataset includes 903,021 individual flights. Our final winter dataset includes 197,081 individual flights.

IV.C. Measurement Issues and Variables

---

[14] We exclude routes that have either endpoint in Alaska, Hawaii, Puerto Rico, Guam or the U.S. Virgin Islands.

[15] These are airports for which either the weather data are missing frequently, or which were entered or exited during the sample period.

We now describe our variable construction. Variable names and definitions appear in Table 2a. Summary statistics are in Table 2b.

i. Dependent Variables

We construct two dependent variables that we use throughout the empirical analysis. The first, **Departure Delay,** measures the time between the scheduled departure and the actual departure of the aircraft from the gate.[16] If the actual departure takes place before the scheduled departure (i.e.: a flight departs early), we set **Departure Delay** to zero.[17] As reported in Table 2b, the average delay in our summer sample is 13.7 minutes while the average delay in our winter sample is 14.0 minutes.

Our second dependent variable is a dummy variable, **Cancelled**, that equals one if the flight is cancelled. As Table 2b shows, on average, 4% of flights are cancelled in our summer sample and 7% of flights are cancelled in our winter sample.

We choose not to include arrival delays (i.e. the time between scheduled arrival and actual arrival) in our analysis because arrival delays are influenced by wind conditions during the flight, and are thus a fairly noisy measure of an airline's performance. In contrast, both departure delays and cancellations are to a larger extent under the control of the airline.

ii. Ownership and Scale of Operations Variables

To measure the extent of a major's vertical integration with its regionals at an airport, we construct **Fraction Owned** which measures the fraction of all regional flights that a major has departing from an airport that are operated by an owned regional partner.[18] As Table 2b indicates, the mean of **Fraction Owned Regional** is 0.56 in the summer sample and 0.41 in the winter sample.

Airlines with larger total operations at an airport may be differentially affected by adverse weather at that airport. Therefore, we control for the size of a major's total, as

---

[16] Thus, our delay measure does not include delays that occur on the runway. We do this intentionally since delays on the runway are less likely to be under the airline's control.

[17] We run a robustness check in which we leave early departures as negative delays and the results are robust to this change.

[18] Note that some majors use owned as well as independent regionals at the same airport. **Fraction Owned** can therefore take on other values than 0 and 1.

well as regional, operations at an airport. Using the OAG data, we construct **Total Flights** which equals the total number of flights per day that a major has departing from an airport (including regional flights). We also construct **Regional Flights** which equals the total number of daily regional flights that a major has departing from an airport. Our regression includes interactions of **Total Flights**, **Regional Flights** and **Fraction Owned Regional** with our measures of adverse weather.

iii. Weather Variables

The NOAA data contain daily observations from airport weather stations on the minimum, average and maximum temperature, and the total accumulated precipitation (measured in inches). Based on these data, we construct the following two variables: **Rain** which measures precipitation on days on which the average temperature is above 32 degrees Fahrenheit and **Snow** which measures precipitation on days on which the average temperature is 32 degrees Fahrenheit or less.[19] Note that while **Snow** is not defined for the summer sample (since we exclude any days during this period on which the temperature is below freezing), **Rain** is defined for the winter sample since above freezing temperatures do occur in that sample. As Table 2b shows, the average daily rainfall in the summer sample is 0.10 inches and the average daily rainfall in the winter sample is 0.04 inches. Average daily snowfall in the winter is 0.46 inches.

We define adverse weather in two different ways. First, we use **Rain** and **Snow** directly. Second, we specifically attempt to capture extreme weather. To do this, we calculate the 95[th] percentile of the daily rain distribution during the summer months for each airport in our summer sample. We then construct the dummy variable **Rain>95[th] Percentile** which equals one if the observed rainfall at the departure airport of a route exceeds the 95[th] percentile of that airport's rain distribution. Note that flights for which this dummy is equal to zero do not have perfect flying conditions. Rather, we are trying to capture the difference between the worst days and all other days with this variable. The mean of **Rain** for observations with **Rain>95[th] Percentile** equal to one is 0.46 inches in the summer sample. We define the variables **Rain>95[th] Percentile** and **Snow>95[th]**

---

[19] We assume an average water equivalent for snow of 8%, i.e. we convert 0.01 inch of accumulated precipitation on days with below freezing temperatures into 0.125 inches of accumulated snow.

Percentile analogously for the rain and snow distributions in the winter sample. The mean of ***Rain*** for observations with ***Rain>95th Percentile*** equal to one is 0.57 inches in the winter sample, and the mean of ***Snow*** for observations with ***Snow>95th Percentile*** equal to one is 6.15 inches.

Using the within-airport rain distribution to identify days with "extreme weather" has two benefits. First, it accounts for the fact that the same weather occurrence may have a different impact at different airports, depending on that airport's regular weather patterns. This is particularly important for the snow measure since a small amount of snow will generally be a much bigger disruption in a city that does not usually experience much snow than in a city with regular snowfall. Second, this approach to defining days with extreme weather ensures that bad weather events are observed at all airports in our sample. This, in turn, allows us to exploit the full distribution of owned regional use across all airports in our sample. In contrast, if we defined extreme weather based on an absolute amount of rain or snow, then we would only observe extreme weather events at a small set of airports and we would only be able to exploit variation in owned regional use across that set of airports. Especially because owned regional use will be correlated with average weather patterns, this alternate approach may not provide sufficient variation in owned regional use to carry out our empirical exercise.

iv. Other Controls - Airport Characteristics

We construct several variables that measure airport characteristics that can affect departure delays and/or the likelihood of cancellations. We construct these variables for both the departure and arrival airports of a flight. However, in some specifications, the departure airport variables will not be separately identified from the fixed effects that we include. Note that conditions at the arrival airport can affect departure delays, especially if the arrival airport has issued a so-called ground stop, which orders all flights that are scheduled for landing to remain at their departure airport until the ground stop is lifted. We do not include any airline-specific characteristics of the arrival airport because ground stops are general and not specific to any particular carrier.

The first variable that we construct is **Slot** which is a dummy for whether the departure (arrival) airport is slot-controlled.[20] We expect delays to be greater at slot-controlled airports. We further control for **Airport Flights**, the total number of domestic flights scheduled to depart from (arrive at) an airport on a given day.[21] This provides an additional measure of airport size. For both the departure and the arrival airport, we interact the **Airport Flights** variable with the variables measuring weather on that day. Since the main effect of bad weather is to require greater time between takeoffs and landings, the effect of adverse weather should be greater at more congested airports.

v. Instruments

Recall that the logic of our instruments is that characteristics of the endpoint airports of routes served by a regional partner from a given airport will be correlated with the type of regional that the major chooses to use for those routes. Based on FL, the four endpoint characteristics that we measure are: whether the route arrives at the major's hub, the average annual precipitation at the endpoint airport, the average annual snowfall at the endpoint airport, and the number of months with an average temperature that is below freezing.[22] We take the average of each of these characteristics across all of the routes that major serves from a given airport with a regional partner and use the four resulting variables as our instruments.

V. Results

V.A. First-stage Results

Table 3 presents the results of our first stage regression of **Fraction Owned** on the instruments described in Section IV and on our exogenous variables. We present our first stages using **Rain>95th Percentile** and **Snow>95th Percentile** as our weather measures. The results are quite similar when we use **Rain** and **Snow** instead.

---

[20] In our sample, the slot-controlled airports are Chicago O'Hare, John F. Kennedy in New York, and Reagan National in Washington, DC. We have excluded LaGuardia Airport in New York (see above).

[21] **Airport Flights** is constructed from the OAG data which only provide a representative flight schedule for one week of each quarter. Therefore, **Airport Flights** takes the same value for each Monday of a quarter, each Tuesday of a quarter, etc…

[22] We construct these instruments using the same data used there.

We show the results for the summer sample in column (3-1) and the results for the winter sample in column (3-2). Robust standard errors are in parentheses. Recall that our instruments are the average characteristics of the arrival airports served by a major's regional partners from the departure airport of the major's route. All of the instruments have highly significant effects in both samples and the signs of the effects are as in FL. Owned regionals are more likely to be used when more of the regionals' routes arrive at a hub and when the endpoints served by the regionals experience more rain and snow throughout the year. Endpoints with more months with below freezing temperatures are less likely to be served by owned regionals.[23] Joint significance of the instruments is confirmed by the F-statistics presented at the bottom of Table 3. Most of the other explanatory variables also have highly significant coefficients.[24] The R-squared of the regression is 0.57 in both samples.

V.B. Performance Effects – Summer Sample

We now present the results from the estimation of our performance equation. The results from the summer sample appear in Tables 4 and 5 while the results for the winter sample appear in Table 6. In Table 4, we estimate the average effect of **Fraction Owned** on delays and cancellations and in Tables 5 and 6 we add interaction terms between **Fraction Owned** and various weather measures. Because we have a large number of independent variables, each of our tables is divided into several panels. In the first panel, we present **Total Flights** – the major's total number of flights, including its regional flights – and its interaction with our weather measures. The second panel includes **Regional Flights** – the major's total number of regional flights – and its interaction with our weather measures. The third panel includes our measure of vertical integration, **Fraction Owned**, and its interaction with the weather measures. Because this panel includes our primary variables of interest, we highlight it in each of the tables. The final

---

[23] Forbes and Lederman (2007a) explain that this result is consistent with the observation that those airports have shorter delays on average.

[24] The first-stage regressions also include departure airport-date fixed effects and interactions of the instruments with the rain and snow measures. The coefficients on those variables are not reported but are available upon request. The results of the first stage regression for the interactions of **Fraction Owned Regional** with the weather variables are also available upon request.

panels in each table include the departure and arrival airport control variables that are not captured by the fixed effects included in the model.

We begin Table 4 by estimating our performance equation using ordinary least squares (OLS) and without including any fixed effects. While this is not our preferred specification (given the large number of unobservable factors that can affect flight delays), it provides a useful starting point. In particular, it allows the coefficients on all of the control variables, many of which will later be absorbed by fixed effects, to be directly estimated. The standard errors we present are clustered on departure airport-day.

The results for the first two panels of variables are generally consistent across all of our specifications. The coefficient estimates on *Total Flights* and *Regional Flights* indicate that flight delays are decreasing in an airline's total number of flights at the airport, but increasing in the airline's number of regional flights. The first effect suggests that airlines with more total flights at an airport are better able to manage delays.[25] However, controlling for an airline's total number of flights, having more regional flights (of either type) at an airport leads to longer departure delays (on flights operated by the major itself). This second effect likely results from the fact that the small planes operated by regional airlines have slower takeoff and cruising speeds than large jet aircraft. As a result, large jets cannot take off as quickly after small aircraft as they would after other large jets. Since airlines tend to have many of their flights take off at the same time to facilitate passenger connections, this will lead to longer delays for a carrier that has a large number of regional flights departing from an airport.[26]

The interactions of *Total Flights* with the dummy for rain being above the airport's 95[th] percentile are generally insignificant in the delay regressions indicating that an airline's overall size at an airport does not change how it is affected by rain. In contrast, the coefficients on the interactions of *Regional Flights* with *Rain>95[th] Percentile* are positive and significant suggesting that airlines with more regional flights experience longer delays on days with rain above the airport's 95[th] percentile of the rain distribution. This is consistent with the direct effect of *Regional Flights*.

---

[25] Note that we separately control for the total number of flights at the airport by all airlines.
[26] This explanation is consistent with Rupp (2005) which finds that smaller aircraft experience significantly longer flight delays.

We now turn to our main variable of interest, **Fraction Owned**. Consistent with our hypothesis, the direct effect of this variable is negative and statistically significant, indicating that majors that use owned regionals for a larger fraction of their regional flights at an airport experience shorter delays on their own flights departing from that airport. The point estimate in this OLS regression implies that majors using only owned regionals experience delays that are 5.4 minutes shorter on average than those experienced by majors using only independent regionals.

The departure and arrival airport control variables have the expected signs. Flights departing from or arriving at slot-controlled airports and airports with more total flights experience longer departure delays. Flights departing from or arriving at airports on days with rainfall above the airport's 95^{th} percentile experience significantly longer departure delays – approximately 13 minutes if the rain is at the departure airport and 10 minutes if the rain is at the arrival airport. The positive coefficients on the interactions of ***Rain>95th Percentile*** with the airports' number of flights indicate that the impact of rain on delays increases with airport congestion. This is precisely what one would expect given that the main effect of rain is to require more time between takeoffs and landings.

In the second column of the Table 4, we re-estimate (4-1) with two-staged least squares (2SLS), instrumenting for **Fraction Owned**. The coefficient estimates on the exogenous variables are almost identical to the OLS results in (4-1). The coefficient on **Fraction Owned** increases slightly in magnitude suggesting an upward bias in the OLS estimate. The 2SLS estimate implies that majors using only owned regionals experience delays that are 6.1 minutes shorter than those experienced by majors using only independent regionals.

In the third column of the table, we add departure airport fixed effects to the model. These control for average differences in departure delays across airports in our sample. The **Slot** and **Airport Flights** variables for the departure airport are not separately identified from these fixed effects. The inclusion of the fixed effects slightly reduces the estimate on **Fraction Owned**. The coefficients on all of the other variables are virtually unchanged.

In the fourth column of the table, we include departure airport-day fixed effects. By doing so, we are able to control for the average delay at every airport in our sample on

each day. This is important for several reasons. First, there are a number of factors that will affect delays at an airport on a given day that are unobservable to us, for example, air traffic control disruptions or security breaches. To the extent that these affect all airlines at an airport equally, they will be captured by the fixed effects. Second, our weather data do not include some elements, such as fog and cloud cover, that affect flight delays and cancellations; these too will be captured by the fixed effects. Finally, the relationship between weather characteristics and delays and cancellations may vary with a host of airport characteristics. Departure airport-day fixed effects allow us to control for this relationship in the most flexible way possible. Note that with the inclusion of these fixed effects, none of the departure airport control variables are separately identified.

Our regression results are quite similar when we include these fixed effects. We present robust standard errors for this and all following specifications. The estimate on *Fraction Owned* implies that majors using only owned regionals at an airport experience flight delays that are approximately 5.6 minutes shorter than the delays experienced by majors *at the same airport and on the same day* using only independent regionals. This effect is quite substantial when compared to the average departure delay in this sample which is 13.7 minutes.

The final column of Table 4 re-estimates (4-4) using *Cancelled* as the dependent variable.[27] The estimate on *Fraction Owned* is again negative and significant indicating that majors using owned regionals not only have shorter delays but are also less likely to have flight cancellations. The point estimate suggests that using only owned regionals lowers the likelihood of a major's flight being delayed by about 1.3 percentage points, compared to majors using only independent regionals. This is a large effect given that the in the summer sample only about 3.7 percent of flights are cancelled.

The results in the first two panels of (4-5) are somewhat different from the results on these variables in the delay regressions. In particular, whereas majors with more total flights at an airport experience shorter delays on average, they do not experience fewer cancellations. Similarly, while majors with more regional flights experience longer delays, there is no effect of *Regional Flights* on cancellations. The results in (4-5)

---

[27] Note that the number of observations varies between the specifications with *Departure Delay* and *Cancelled* because *Departure Delay* is missing for flights that are cancelled.

further indicate that on days with rainfall above the airport's 95[th] percentile, carriers with more total flights at the airport experience more cancellations but carriers with more regional flights experience fewer cancellations. Since having more total flights is positively correlated with an airline's average flight frequency on a route but, conditional on the carrier's total number of flights, having more regional flights is negatively correlated with flight frequency (since regionals are typically used on low-demand routes), these results are consistent with Rupp (2005) which finds that cancellations are more common on routes with high flight frequencies. Rupp also finds that, controlling for flight frequency, smaller planes have longer delays than larger planes but are no more likely to be cancelled.

In Table 5, we add interactions between *Fraction Owned* and various measures of rain. Recall that our hypothesis is that majors with owned regionals should not only have an overall performance advantage over majors with independent regionals, but this performance advantage should also increase as weather deteriorates because adverse weather increases the need for unforeseen schedule adjustments. The specifications in this table allow us to test both parts of this hypothesis. We include departure airport-date fixed effects in all specifications presented in this table. We find that the estimates on all of the exogenous variables are quite similar to Table 4. We therefore focus our discussion on the results we find for the vertical integration measures.

In the first column of Table 5 we present OLS estimates, and in the remaining columns we present 2SLS estimates instrumenting for both *Fraction Owned* and its interaction with the rain measures. As in Table 4, the estimate on *Fraction Owned* in (5-1) is negative and statistically significant. The estimate on *Rain>95[th] Percentile* interacted with *Fraction Owned* is also negative and statistically significant. Consistent with our hypothesis, this suggests that using owned regionals provides majors with a performance advantage on all days and that this advantage increases as the weather deteriorates. Note that the uninteracted *Fraction Owned* variable should not be interpreted as measuring the performance advantage of owned regionals on days with ideal flying conditions (or days with zero schedule disruptions). Rather, since *Rain>95[th] Percentile* only captures the very rainiest days at each airport, the uninteracted *Fraction Owned* term includes all other days. Many of these days will include schedule

25

disruptions that occur for weather-related or other reasons. Thus, it is not surprising that we find a performance advantage on these days as well.

In the second column, we instrument for the ownership variable and its interaction. As in Table 4, the 2SLS estimate on the direct effect of **Fraction Owned** is barely different from the OLS estimate. However, the 2SLS estimate on the interaction term more than doubles in magnitude. This suggests that endogeneity of the vertical integration decision is indeed a concern when estimating its effect on performance. In particular, the upward bias of the coefficient in the OLS regression is consistent with the implication from Figure 1 – namely, that simply comparing the performance of integrated and non-integrated firms will wrongly suggest that integrated firms perform worse, because firms choose to integrate more difficult transactions. The instrumental variables approach accounts for this endogeneity of the ownership decision.

The results from the 2SLS regression confirm our hypothesis that carriers who are vertically integrated with their regionals perform better than carriers that are vertically separated and that the performance advantage of integrated carriers is greater in bad weather conditions. The coefficient on the direct effect of **Fraction Owned** implies that - on days with rainfall *below* the 95th percentile of the airport's distribution -   majors using only owned regionals at the airport experience flight delays that are approximately 5 minutes shorter than the delays experienced by majors using only independent regionals. The coefficient on the interaction term implies that - on days with rain *above* the 95th percentile - the performance advantage of majors who only use owned regionals increases so that their delays are 13 minutes shorter than the delays of majors who only use independent regionals. It is important to note that this specification only provides estimates of the performance *advantage* of owned regionals in "good" and "bad" weather since the direct effect of weather is absorbed by the airport-day fixed effects we include. Comparing our results to the average delay in the sample indicates that the effects we find are again quite substantial. While the overall mean of **Departure Delay** is 13.7 minutes, the mean delay on days with rain below the 95th percentile is 13 minutes while the mean delay on days with rain above the 95th percentile is 28 minutes.

(5-3) presents the same specification but adds a second rain measure. We add the variable **Rain>75th Percentile** which equals one if the rainfall on a given day is above the

26

75th percentile of the rain distribution at the airport in the summer sample. We add this variable to investigate how the performance advantage of owned regionals varies across different parts of the rain distribution. The estimates in (5-3) indicate that the performance advantage of owned regionals does not appear to increase with smaller amounts of rain. While the coefficient on **Rain>75th Percentile*Fraction Owned** is negative, it is not statistically significant. The coefficient on **Rain>95th Percentile*Fraction Owned** is again negative and significant. Its magnitude and significance are both slightly reduced because some of the effect of **Rain>95th Percentile** is now captured by **Rain>75th Percentile*Fraction Owned** since these variables are additive rather than mutually exclusive.

In column four, we further investigate the relationship between the performance advantage of ownership and rain by using the linear **Rain** variable. We use this alternate measure of "bad" weather because it provides an absolute – rather than relative - measure of "bad" weather. It also exploits the full distribution of observed rain. However, based on the results in (5-3), we do not expect the relationship between rain and the performance advantage of owned regionals to be the same at all points of the rain distribution. The coefficient on **Rain*Fraction Owned** is negative but it is statistically insignificant. We interpret the findings in specifications (5-2) through (5-3) as evidence that the effect of rain on the performance advantage of owned regionals is nonlinear.

In the remaining three columns of Table 5, we re-estimate (5-2) through (5-4) using **Cancelled** as the dependent variable. Recall from Table 4 that we found that the use of owned regionals both lowered delays and reduced cancellations. We now investigate whether the relationship between ownership and cancellations is affected by the presence of bad weather. The results in columns five through seven indicate that while majors using owned regionals experience fewer cancellations in any weather, this effect is no larger in "bad" weather. The coefficients on the interaction terms between **Fraction Owned** and all of our rain measures are statistically insignificant. Thus, these results suggest that the performance advantage of owned regionals with respect to cancellations does not increase as weather deteriorates. However, we are cautious to attach too much weight to these findings. Specifically, the inclusion of the departure airport-day fixed effects means that these interaction effects are identified only by

variation across airlines at an airport on a given "bad" weather day. In order for us to estimate whether heavy rain impacts the relationship between owned regional use and cancellations, the decision to cancel flights must be, at least to some extent, under the airline's control. However, cancellations in heavy rain tend to occur because of the presence of heavy storms. In these situations, airlines typically do not have much control over their number of cancellations.[28]

V.C. Performance Effects – Winter Sample

We now turn to our results for the winter sample. In this sample, we include both snow and rain as weather measures; however, we focus on the snow results because rain is observed infrequently and in small amounts during this period. Recall that this sample is restricted to the Midwest and Northeast census regions (i.e. to those geographic regions where all airports experience snow during the winter) during the months December through February, and that we exclude weekend days. The **Snow>95th Percentile** variable therefore captures the three snowiest days at each airport. The results for this sample are presented in Table 6.

In the first column, we present a two-staged least squares specification analogous to that in (5-2) but with the additional snow variables. We instrument for **Fraction Owned**, **Snow>95th Percentile*Fraction Owned** and **Rain>95th Percentile*Fraction Owned**. Consistent with the results from the summer sample, we find a negative and significant direct effect of **Fraction Owned**. The coefficient on the interaction effect with **Snow>95th Percentile** is also negative and is significant at the 10 percent level. The interaction with **Rain>95th Percentile** is insignificant. In fact, we find no statistically significant effects on the interaction of rain with **Fraction Owned** in any of our specifications for the winter sample. This is not surprising since the airports in this sample only experience very few, if any, days with rain during the winter.

Most of the other coefficients in this regression have the same sign as in the summer regressions. For example, we again find that carriers with more flights overall experience shorter delays, whereas carrier with more regional flights of either type

---

[28] Estimation of the cancellation regressions with departure airport fixed effects instead of departure airport-day fixed effects produces negative and statistically significant coefficients on the interaction terms.

28

experience longer delays. With respect to the arrival airport control variables, we again find that flights arriving at airports with more total flights experience longer delays and that this effect is more pronounced when the airport experiences snow or rain above its 95th percentile snow or rain. We find different results from the summer sample on some of the interactions with ***Rain>95th Percentile***. However, again, we do not attach much weight to these interactions because rain is very infrequent in this sample, and some airports experience no rain at all during the winter months.

The coefficient on the direct effect of ***Fraction Owned*** implies that – on days with rain and snow below the 95th percentile of the airport's distribution – majors using only owned regionals experience flight delays that are approximately 1.1 minutes shorter than the delays experienced by majors using only independent regionals. The coefficient on the interaction effect with snow implies that – on days with snow above the 95th percentile – majors using only owned regionals have a performance advantage that translates into delays that are shorter by 8.7 minutes, on average, than for majors using only independent regionals. It is again useful to compare these estimates to the average departure delays in the sample. The mean departure delay in the winter sample on days with snow below the 95th percentile is 14 minutes while the mean delay on days with snow above the 95th percentile is 41 minutes. The magnitude of the performance effects that we estimate in the winter sample are smaller than the ones we find for the summer months; however, given that we are measuring very different types of weather, we do not expect the estimates from the two samples to be identical in size. Rather, we are interested in whether alternate measures of bad weather give qualitatively similar results and the findings so far indicate that they do.

In the second column of Table 6, we present results using the linear measures of ***Snow*** and ***Rain***.[29] As before, this provides us with absolute measures of "bad" weather and exploits the full distribution of rain and snow in the winter sample. We believe that this is a particularly important check in the winter sample because ***Snow>95th Percentile*** captures quite a broad range of snow events. When we use the linear weather variables, we find similar effects on most variables. As in the previous specification, the direct

---

[29] We do not present results with the 75th percentile of the snow or rain distribution here because those percentiles are equal to zero for many airports.

effect of **_Fraction Owned_** and its interaction with **_Snow_** are negative, but now the direct effect is insignificant while the interacted effect is highly significant. The magnitudes imply that an additional inch of snow leads to 2.5 minutes of additional delay for majors which only use independent regionals, compared to majors which only use owned regionals. Interestingly, the results in (6-1) and (6-2) together suggest that – in contrast to the summer results – the performance advantage of owned regionals does not increase only with extreme amounts of precipitation. Rather, in the winter sample, the linear snow measure appears to capture the performance advantage of majors with owned regionals best.

In Columns (6-3) and (6-4) we present results on cancellations. In Column (6-3) we use the weather measures based on the 95th percentiles of the snow and rain distributions, whereas in Column (6-4) we show results for the linear weather measures. We find highly significant negative coefficients on **_Fraction Owned_** and its interaction with the linear snow measure. Interestingly, we find that in the winter sample, carriers with more total flights at an airport experience fewer cancellations while carriers with more total regional flights at an airport experience more cancellations. The results from the summer sample on the relationship between an airline's scale of operations at an airport and cancellations were more ambiguous. This may be because cancellations are three times more likely in the winter sample than in the summer sample which may allow the relationship between scale and cancellations to be more accurately estimated.

The coefficient estimates in (6-4) imply that - on days with no snow or rain - cancellations for majors using only owned regionals at an airport are 1.3 percentage points lower than cancellations for majors using independent regionals. Each additional inch of snow increases the performance advantage of owned regionals by about 1.6 percentage points.

Overall, the results provide strong support for our hypotheses that use of an owned regional allows a major to better optimize its network, in particular when unforeseen schedule adjustments are necessary.

V.D. Robustness

In addition to the specifications reported in Tables 4 through 6, we have also carried out several other robustness checks. Specifically, we have estimated our models using two alternate ways of controlling for a major's overall size of operations and a major's size of regional operations at an airport. This is important because – as FL find – a major's overall scale of operations may be correlated with its decision of what type of regional to use. First, to allow for nonlinear effects, we have included quadratic terms of **Total Flights** and **Regional Flights** (and interact these with our weather variables). When we include these higher order terms, we find that the relationship between delays and these variables are, in fact, nonlinear. In both the summer and winter samples, the coefficients on the **Fraction Owned** terms are qualitatively the same though their magnitudes are slightly reduced. Second, in addition to the continuous **Total Flights** variable, we add a dummy variable that equals one if a flight departs from a major's hub. We also interact the hub variable with our weather variables. The results from both samples are robust to the inclusion of the hub dummy.

As an additional robustness check, we limit our sample to flights that depart from the 50 largest airports in our sample (though not all of these airports appear in the winter sample since it is limited to specific regions). The results are robust to this change in the sample. In addition, we have redefined our dependent variable **Departure Delay** to count early departures as negative delays. Our results are robust to this alternative definition. We have also dropped the 0.5% of our observations with the longest delays to check if our results are sensitive to excluding outliers. Again, we find that our results are robust. Finally, we have added arrival airport fixed effects with no change to our results. Overall, we find strong support for our hypotheses that majors which are vertically integrated with their regionals have shorter flight delays and fewer cancellations in any weather and this performance advantage increases, at least for flight delays, as the weather deteriorates.

VI. Conclusion

In this paper, we have investigated the consequences of vertical integration for the operational performance of major U.S. airlines. These airlines outsource flights to regional airlines, of which some are owned and others are independent and managed

through contractual relationships. Ownership of a regional mitigates incentive problems that arise when unforeseen schedule disruptions occur, causing majors and independent regionals to disagree about how those disruptions should be resolved.

We estimate how specific measures of a major's performance on flights that it operates itself are affected by its degree of vertical integration with regional airlines that operate flights on its behalf out of the same airport. Our empirical investigation takes into account that organizational form decisions are endogenous to firm performance by instrumenting for the extent of a carrier's vertical integration at an airport with the characteristics of the *other airports* served from that airport by the carrier's regional partners. Prior work by Forbes and Lederman (2007a) has shown that these characteristics affect the vertical integration decision, and we find here, as well, that our instruments are highly predictive of vertical integration.

Our empirical approach further takes into account that the relationship between vertical integration and performance may change with the characteristics of a transaction. In our setting, we would expect the performance advantage of owned regionals to increase as weather deteriorates because this causes the likelihood of schedule disruptions to increase. Typically, this type of relationship can be hard to estimate because integration and non-integration are only observed in specific transaction environments. However, because airlines' ownership decisions are fixed in the short-run while weather changes on a daily basis, we are able to observe both owned and independent regionals in various weather environments.

Our results show that ownership of a regional does, indeed, improve airlines' network performance and, even more so, in the presence of adverse weather. The estimates that we obtain are both statistically and economically significant and are robust across a variety of specifications. We interpret our results as indicating that, at least in this setting, contracts between majors and independent regionals cannot replicate the incentive alignment that ownership achieves. We want to emphasize, however, that this paper has only estimated one side of the tradeoff that airlines face when deciding whether to own their regional partners. While our results indicate that there are benefits to ownership in terms of network performance, they do not imply that all majors should

own their regionals. As FL explain, there are costs associated with ownership and ownership is only optimal when the benefits outweigh the costs.

We believe that this paper contributes to the existing literature in several ways. First, we provide one of the few pieces of *direct* evidence that there are indeed performance consequences of vertical integration. We are the first to exploit high frequency data on performance and on variables that affect the returns to vertical integration to trace out how the relationship between vertical integration and performance changes with characteristics of the transaction environment. This allows us to test not only the implication that – if ownership solves incentive problems that arise under incomplete contracts – there will be performance differences on average, but also that these performance differences will increase as the incentive problems grow.

In addition, we exploit specific institutional features of the industry we study to derive instruments for firms' organizational form decisions. Finally, we show that, in a network industry, a firm's vertical integration decision on one set of transactions can have implications for its performance on a different set of transactions. Our findings suggest that firms in network industries should take these performance effects into account when they decide whether to outsource part of their network.

References:

Ahmadjian, Christina and Joanne Oxley (2007), "Vertical Relationships, Hostages, and Supplier Performance: Evidence from the Japanese Automotive Industry", *mimeo*, University of Toronto.

Alchian, Armen A. and Harold Demsetz (1972), "Production, Information Costs, and Economic Organization", *American Economic Review* 62(5), 777-795.

Anderson, Erin and David Schmittlein (1984), "Integration of the Sales Force: An Empirical Examination", *Rand Journal of Economics* 15, 385-95.

Ball, Michael, Cynthia Barnhart, George Nemhauser and Amedeo Odoni (2005), "Managing Air Traffic and Airline Operations for Schedule Reliability", *mimeo*, Massachusetts Institute of Technology.

Ciliberto, Frederic (2006), "Does Organizational Form Affect Investment Decisions?", Journal of Industrial Economics 54 (1), 63-93.

Coase, Ronald (1937), "The Nature of the Firm", *Economica* 4, 386-405.

Forbes, Silke J. (2007), "The Effect of Air Traffic Delays on Airline Prices", *mimeo*, University of California, San Diego

Forbes, Silke J. and Mara Lederman (2007a), "Control Rights, Network Structure and Vertical Integration: Evidence from Regional Airlines", *mimeo*, University of California, San Diego, and University of Toronto

Forbes, Silke J. and Mara Lederman (2007b), "The Role of Regional Airlines in the U.S. Airline Industry", *mimeo*, University of California, San Diego, and University of Toronto

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.

Holmström, Bengt (1982), "Moral Hazard in Teams", *Bell Journal of Economics* 13, 324-340.

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 and Margaret Slade (2007), "Vertical Integration and Firm Boundaries: The Evidence", *mimeo*, University of Michigan.

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

Masten, Scott (1993), "Transaction Costs, Mistakes, and Performance: Assessing the Importance of Governance", *Managerial and Decision Economics* 14(2), 119-129.

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.

Masten, Scott E., James W. Meehan, Jr., and Edward A. Snyder (1991), "The Costs of Organization", *Journal of Law, Economics, and Organization* 7(1), 1-25.

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

Novak, Sharon and Scott Stern (forthcoming), "How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry", *Management Science*

Rupp, Nicholas (2005), "Flight Delays and Cancellations", *mimeo*, East Carolina University.

Rupp, Nicholas and George M. Holmes (forthcoming), "An Investigation Into the Determinants of Flight Cancellations", *Economica*.

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

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

Williamson, Oliver E. (1991), "Comparative Economic Organization: The Analysis of Discrete Structural Alternatives", *Administrative Science Quarterly* 36(2), 269-296.

Table 1
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/CCAir
	Chautauqua Airlines
	Colgan Airways
	Commutair
	Mesa Air Group/Mesa Airlines
	Piedmont Airlines
	PSA Airlines

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

Table 2a
Variable Names and Definitions

Variable	Definition	Source
DEPENDENT VARIABLES		
Departure Delay	Difference between scheduled departure and actual departure of aircraft from the gate; =0 if actual departure is before scheduled departure	BTS On-time data
Cancelled	=1 if flight is cancelled	BTS On-time data
OWNERSHIP VARIABLE		
Fraction Owned Regional	Fraction of major's regional flights at the departure airport that are operated by an owned regional partner	OAG & RAA data
SCALE OF OPERATIONS VARIABLES		
Total Flights	A carrier's total number of flights at the departure airport on a day (including regional flights), in hundreds	OAG data
Regional Flights	A carrier's total number of *regional* flights at the departure airport on a day, in hundreds	OAG data
WEATHER VARIABLES (defined for both departure and arrival airports)		
Rain	Daily precipitation, on days with average temperature >32 degrees Fahrenheit (inches)	NOAA data
Rain>75th Percentile	=1 if rain at an airport on a day is greater than the 75th percentile rain observed at that airport during the summer (winter) sample	NOAA data
Rain>95th Percentile	=1 if rain at an airport on a day is greater than the 95th percentile rain observed at that airport during the summer (winter) sample	NOAA data
Snow	Daily precipitation, on days with average temperature <=32 degrees Fahrenheit (inches)	NOAA data
Snow>95th Percentile	=1 if snow at an airport on a day is greater than the 95th percentile snow observed at that airport during the winter sample	NOAA data
OTHER CONTROLS (defined for both departure and arrival airports)		
Airport Flights	Total number of domestic flights scheduled to depart from (arrive at) the airport on a day, in hundreds	OAG data
Slot	=1 if the airport is one of four slot-controlled airports (ORD, LGA, JFK, DCA)	Authors' construction

Table 2b
Summary Statistics

	Summer Sample				Winter Sample			
	Mean	St Dev	Min	Max	Mean	St Dev	Min	Max
DEPENDENT VARIABLES								
Departure Delay (min)	13.7	35.3	0	1435	14.5	33.6	0	1129
Cancelled	0.04	0.19	0	1	0.07	0.25	0	1
OWNERSHIP VARIABLE								
Fraction Owned Regional	0.56	0.45	0	1	0.41	0.45	0	1
SCALE OF OPERATIONS VARIABLES								
Total Flights (in hundreds)	3.45	2.71	0.03	9.26	3.07	1.99	0.03	5.86
Regional Flights (in hundreds)	1.06	0.80	0.01	2.98	0.91	0.56	0.01	1.98
WEATHER VARIABLES (departure airports)								
Rain (inches)	0.10	0.35	0	12.56	0.04	0.16	0	2.24
*Rain	Rain>75th Percentile=1*	0.46	0.63	0.01	12.56				
*Rain	Rain>95th Percentile=1*	1.21	1.025	0.02	12.56	0.58	0.45	0.01	2.24
Snow (inches)					0.46	1.57	0	17.13
Snow/ Snow>95th Percentile=1					6.36	3.29	1.4	17.13
OTHER CONTROLS (departure airports)								
Airport Flights (in hundreds)	6.36	3.42	0.14	12.38	6.08	3.07	0.38	11.98
Slot	0.13	0.33	0	1	0.22	0.42	0	1

Table 3
First Stage Regression

Dependent Variable	*Fraction Owned Regional*	
Sample	Summer	Winter
	(3-1)	(3-2)
INSTRUMENTS		
Fraction of Regional's Routes Arriving at Hub	0.2350	0.2196
	(0.0024)**	(0.0066)**
Average Annual Precipitation at Endpoints Served by Regional	0.0039	0.0192
	(0.0001)**	(0.0004)**
Average Annual Snowfall at Endpoints Served by Regional	0.0095	-0.4327
	(0.0001)**	(0.0065)**
Average # of Months with Below Freezing Temperature at Endpoints Served by Regional	-0.2424	0.0160
	(0.0018)**	(0.0002)**
TOTAL # OF FLIGHTS		
Total Flights	-0.1373	-0.3006
	(0.0013)**	(0.0020)**
*Rain>95th Percentile*Total Flights*	0.0780	0.0523
	(0.0062)**	(0.0089)**
*Snow>p95*Total Flights*		-0.0492
		(0.0101)**
TOTAL # OF REGIONAL FLIGHTS		
Regional Flights	0.5048	1.1280
	(0.0039)**	(0.0041)**
*Rain>95th Percentile*Regional Flights*	-0.2656	-0.0344
	(0.0183)**	(0.0182)+
*Snow>p95*Regional Flights*		-0.0507
		(0.0231)*
ARRIVAL AIRPORT CONTROLS		
Slot	-0.0580	-0.0487
	(0.0013)**	(0.0020)**
Airport Flights	0.0069	0.0036
	(0.0001)**	(0.0002)**
Rain>95th Percentile	-0.0042	-0.0013
	(0.0017)*	(0.0030)
Snow>p95		0.0027
		(0.0039)
*Rain>95th Percentile*Airport Flights*	-0.0030	-0.0000
	(0.0005)**	(0.0009)
*Snow>p95*Airport Flights*		-0.0009
		(0.0011)
Observations	869,401	183,453
F-statistic on instruments	$F_{(4,860,428)} = 8493.20$	$F_{(4,183,428)} = 4994.35$
Prob>F	0.0000	0.0000

R-squared	0.57	0.57

Robust standard errors in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. Specifications include departure airport-date fixed effects and interactions of the instruments with the weather variables (coefficients not reported).

Table 4
Effect of Owned Regional Use on Delays and Cancellations
Direct Effect Only
Summer Months

Dependent Variable	*Departure Delay (min)*				*Cancelled*
Estimation Method	OLS	2SLS	2SLS	2SLS	2SLS
Fixed Effects	None	None	Departure Airport	Departure Airport-Date	Departure Airport-Date
	(4-1)	(4-2)	(4-3)	(4-4)	(4-5)
TOTAL # OF FLIGHTS					
Total Flights	-0.90	-0.97	-0.25	-0.75	-0.0001
	(0.15)**	(0.15)**	(0.21)	(0.11)**	(0.0006)
Rain>95th Percentile*Total Flights	-2.00	-1.97	-1.82	-0.59	0.0126
	(1.62)	(1.62)	(1.60)	(0.66)	(0.0033)**
TOTAL # OF REGIONAL FLIGHTS					
Regional Flights	2.29	2.56	1.81	3.66	-0.0005
	(0.50)**	(0.50)**	(0.69)**	(0.34)**	(0.0018)
Rain>95th Percentile*Regional Flights	9.64	9.54	9.16	5.82	-0.0319
	(5.03)+	(5.02)+	(4.97)+	(2.02)**	(0.0097)**
FRACTION OWNED REGIONAL					
Fraction Owned	-5.40	-6.14	-5.50	-5.60	-0.0120
	(0.32)**	(0.58)**	(0.63)**	(0.48)**	(0.0022)**
DEPARTURE AIRPORT CONTROLS					
Slot	2.48	2.37			
	(0.89)**	(0.88)**			
Airport Flights	0.59	0.61			
	(0.07)**	(0.07)**			
Rain>95th Percentile	12.54	12.55	12.40		
	(3.11)**	(3.11)**	(3.08)**		
Rain>95th Percentile*Airport Flights	1.52	1.51	1.44		
	(0.79)+	(0.79)+	(0.78)+		
ARRIVAL AIRPORT CONTROLS					
Slot	1.37	1.29	0.86	0.96	0.0194
	(0.19)**	(0.19)**	(0.19)**	(0.14)**	(0.0008)**
Airport Flights	0.13	0.14	0.13	0.13	0.0013
	(0.02)**	(0.02)**	(0.02)**	(0.01)**	(0.0001)**
Rain>95th Percentile	10.25	10.24	10.20	8.15	0.0410
	(0.44)**	(0.44)**	(0.44)**	(0.27)**	(0.0014)**
Rain>95th Percentile*Airport Flights	1.78	1.78	1.77	1.82	0.0089
	(0.13)**	(0.13)**	(0.13)**	(0.09)**	(0.0004)**
Observations	869,401	869,401	869,401	869,401	903,021

Robust standard errors in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. In (4-1)-(4-3), standard errors are clustered on departure airport-date.

Dependent Variable	*Departure Delay (min)*				*Cancelled*		
Estimation Method	OLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Fixed Effects	Departure Airport-Date						
	(5-1)	(5-2)	(5-3)	(5-4)	(5-5)	(5-6)	(5-7)
TOTAL # OF FLIGHTS							
Total Flights	-0.68	-0.70	-0.54	-0.74	0.0001	-0.0006	-0.0004
	(0.08)**	(0.11)**	(0.11)**	(0.11)**	(0.0006)	(0.0006)	(0.0006)
*Rain>75th Percentile*Total Flights*			-0.64			0.0032	
			(0.31)*			(0.0016)*	
*Rain>95th Percentile*Total Flights*	-0.92	-1.40	-0.93		0.0097	0.0072	
	(0.66)	(0.74)+	(0.79)		(0.0038)**	(0.0040)+	
*Rain*Total Flights*				-0.32			0.0068
				(0.49)			(0.0023)**
TOTAL # OF REGIONAL FLIGHTS							
Regional Flights	3.43	3.48	2.91	3.59	-0.0010	0.0012	0.0006
	(0.24)**	(0.34)**	(0.35)**	(0.35)**	(0.0018)	(0.0019)	(0.0018)
*Rain>75th Percentile*Regional Flights*			2.79			-0.0085	
			(1.00)**			(0.0051)+	
*Rain>95th Percentile*Regional Flights*	6.82	8.25	6.18		-0.0234	-0.0169	
	(2.00)**	(2.28)**	(2.44)*		(0.0113)*	(0.0121)	
*Rain*Regional Flights*				4.03			-0.0148
				(1.43)**			(0.0067)*
FRACTION OWNED REGIONAL							
Fraction Owned	-5.15	-5.25	-4.87	-5.56	-0.0111	-0.0101	-0.0139
	(0.12)**	(0.48)**	(0.52)**	(0.51)**	(0.0022)**	(0.0024)**	(0.0023)**
*Rain>75th Percentile*Fraction Owned*			-1.45			-0.0039	
			(1.24)			(0.0058)	
Rain>95th Percentile Fraction Owned*	-3.05	-7.97	-6.23		-0.0217	-0.0163	
	(0.90)**	(3.34)*	(3.49)+		(0.0149)	(0.0157)	
Rain Fraction Owned*				-1.92			0.0029
				(1.77)			(0.0065)
ARRIVAL AIRPORT CONTROLS							
Slot	0.99	0.96	0.78	0.71	0.0194	0.0186	0.0180
	(0.14)**	(0.14)**	(0.14)**	(0.14)**	(0.0008)**	(0.0008)**	(0.0008)**
Airport Flights	0.12	0.13	-0.01	0.07	0.0013	0.0009	0.0011
	(0.01)**	(0.01)**	(0.01)	(0.01)**	(0.0001)**	(0.0001)**	(0.0001)**
Rain>75th Percentile			3.50			0.0170	
			(0.12)**			(0.0006)**	
Rain>95th Percentile	8.14	8.13	5.34		0.0410	0.0275	
	(0.27)**	(0.27)**	(0.29)**		(0.0014)**	(0.0015)**	

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Rain				7.05 (0.21)**			0.0376 (0.0011)**
Rain>75th Percentile*Airport Flights			0.81 (0.03)**			0.0022 (0.0002)**	
Rain>95th Percentile*Airport Flights	1.82 (0.09)**	1.82 (0.09)**	1.16 (0.09)**		0.0089 (0.0004)**	0.0071 (0.0005)**	
Rain*Airport Flights				1.88 (0.07)**			0.0080 (0.0003)**
Observations	869,401	869,401	869,401	869,401	903,021	903,021	903,021

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

Table 6
Effect of Owned Regional Use on Delays and Cancellations
Direct and Interaction Effects
Winter Months

Dependent Variable	*Departure Delay (min)*		*Cancelled*	
Estimation Method	2SLS	2SLS	2SLS	2SLS
	(6-1)	(6-2)	(6-3)	(6-4)
TOTAL # OF FLIGHTS				
Total Flights	-0.68	-0.34	-0.0177	-0.0140
	(0.27)*	(0.29)	(0.0019)**	(0.0020)**
Snow>95th Percentile*Total Flights	-3.67		-0.0275	
	(2.93)		(0.0177)	
Rain>95th Percentile*Total Flights	2.18		0.0136	
	(1.18)+		(0.0090)	
Snow*Total Flights		-1.34		-0.0089
		(0.54)*		(0.0026)**
Rain*Total Flights		3.86		-0.0027
		(1.89)*		(0.0119)
TOTAL # OF REGIONAL FLIGHTS				
Regional Flights	3.83	2.53	0.0461	0.0334
	(0.83)**	(0.89)**	(0.0057)**	(0.0059)**
Snow>95th Percentile*Regional Flights	15.64		0.1272	
	(8.22)+		(0.0500)*	
Rain>95th Percentile*Regional Flights	-4.26		-0.0569	
	(3.47)		(0.0266)*	
Snow*Regional Flights		4.88		0.0328
		(1.56)**		(0.0080)**
Rain*Regional Flights		-8.01		-0.0234
		(5.83)		(0.0392)
FRACTION OWNED REGIONAL				
Fraction Owned	-1.08	-0.52	-0.0192	-0.0130
	(0.59)+	(0.61)	(0.0037)**	(0.0038)**
Snow>95th Percentile*Fraction Owned	-8.58		-0.0387	
	(4.83)+		(0.0330)	
Rain>95th Percentile*Fraction Owned	0.62		0.0091	
	(2.56)		(0.0193)	
Snow*Fraction Owned		-2.53		-0.0158
		(0.92)**		(0.0052)**
Rain*Fraction Owned		0.34		0.0056
		(3.90)		(0.0275)
ARRIVAL AIRPORT CONTROLS				
Slot	-0.35	-0.46	0.0296	0.0219
	(0.24)	(0.24)+	(0.0021)**	(0.0020)**

Airport Flights	0.15 (0.02)**	0.15 (0.02)**	0.0021 (0.0002)**	0.0022 (0.0002)**
Snow>95th Percentile				
Rain>95th Percentile	3.06 (0.43)**		0.0327 (0.0033)**	
Snow		1.06 (0.14)**		0.0395 (0.0008)**
Rain		3.98 (0.50)**		0.0489 (0.0035)**
Snow>95th Percentile*Airport Flights	1.08 (0.20)**		0.0140 (0.0014)**	
Rain>95th Percentile*Airport Flights	0.77 (0.12)**		0.0037 (0.0010)**	
Snow*Airport Flights		0.24 (0.04)**		0.0020 (0.0002)**
Rain*Airport Flights		0.39 (0.12)**		0.0044 (0.0010)**
Observations	183,453	183,453	197,081	197,081

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