

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

Silke J. Forbes
University of California, San Diego

Mara Lederman
Rotman School of Management, University of Toronto

This draft: April 2010

Abstract

This paper investigates the effects of vertical integration on operational performance in the U.S. airline industry. All of the large U.S. network carriers use regional partners to operate some of their short- and medium-haul routes. However, some regional partners are owned while others are independent and managed through contracts. Using detailed flight-level data and accounting for the potential endogeneity of integration decisions, we estimate whether an airline's use of an owned - rather than independent - regional partner at an airport affects its delays and cancellations on the flights that it itself operates out of that airport. Our results indicate that integrated airlines perform systematically better than non-integrated airlines *at the same airport on the same day*. Furthermore, we find that the performance advantage of integrated airlines increases on days with adverse weather conditions and when airports are more congested. These findings suggest that, in this setting, vertical integration may facilitate real-time adaptation decisions. We believe that this work is one of the first to both document the performance implications of vertical integration decisions as well as shed light on the underlying mechanism.

JEL codes: L22, L25

Key words: Vertical Integration; Firm Boundaries; Firm Performance; Airlines

We thank the editor, Ariel Pakes, and two anonymous referees, Itai Ater, Severin Borenstein, Kenneth Corts, Julie Cullen, Luis Garicano, Robert Gibbons, Scott Masten, Anita McGahan, Michael Salinger, Steven Tadelis and seminar/conference participants at the Harvard-MIT Organizational Economics Seminar, Northwestern University (Kellogg School of Management), Loyola Marymount University, the University of California, Berkeley (Haas School of Business), the University of Rochester (Simon School of Business), the University of Southern California (Marshall School of Business), the University of Toronto (Rotman School of Management), the Allied Social Science Association Conference, the International Society for New Institutional Economics Conference and the NBER Universities Research Conference on the Airline Industry for helpful comments and suggestions. Graton Gathright provided excellent research assistance. All errors are our own.

I. Introduction

Do firm boundary decisions affect firm performance? While theoretical work predicts that one should find performance implications of vertical integration decisions, there is little direct evidence on this question.¹ Such empirical evidence has been hard to establish for two reasons. First, it is difficult to obtain data on outcome measures for similar transactions that are organized differently. Second, firm boundary decisions will typically be endogenous (Masten, 1993). In this paper, we overcome these difficulties and document the existence and magnitude of performance differences between integrated and non-integrated firms carrying out virtually identical transactions. Our results indicate that, operationally, there is a performance advantage to vertical integration.² Moreover, we find that this performance advantage increases when the need for adaptation decisions is greater. We believe that this paper is one of the first to both measure the performance implications of integration decisions as well as provide empirical evidence on a possible underlying cause.

Our setting is the U.S. airline industry. All of the large U.S. network carriers, often called “majors”, employ regional airlines to operate a subset of their routes. Regional airlines may be owned by the major for which they operate or they may be independent and contract with one or more major carriers. To the extent that contracts do not fully align an independent regional’s incentives with those of its major, vertically integrated airlines may perform systematically differently than vertically separated ones. In particular, if independent regionals are less willing than owned regionals to carry out non-contractible adaptation decisions, then ownership may affect the efficiency with which majors can respond to schedule disruptions (which arise frequently in this industry). Moreover, because a regional’s flights and a major’s own flights compete for potentially scarce airline and airport resources, this may impact not only the performance of the regional’s flights but the performance of the major’s as well.

We estimate whether use of an owned – rather than independent – regional at a particular airport affects a major’s performance on the flights that it itself operates out of that airport. We measure performance using flight-level data on delays and cancellations.

¹ See, for example, Williamson (1975, 1985), Grossman and Hart (1986) and Hart and Moore (1990).

² We use the term “operational performance” to distinguish what we measure from measures of overall performance such as profits.

Our empirical approach exploits three important features of our setting. First, because we observe many airports at which some majors use owned regionals while others use independent regionals, we are able to include fixed effects for each origin airport-day combination in our regressions, in addition to a rich set of control variables. Thus, we identify the effects of integration by comparing the performance metrics of integrated and non-integrated majors *at the same airport on the same day*. Second, because we are estimating the impact of ownership of a regional on the performance of the major, characteristics that affect the returns to integration on regional routes but which do not directly affect performance on the major's routes can serve as instruments for airlines' ownership decisions. Finally, weather conditions provide an observable and exogenous source of variation in the likelihood that adaptation decisions will be needed. Because weather changes on a daily basis while ownership decisions are fixed at least in the medium term, we are able to trace out the relationship between integration and performance as the need for adaptation decisions increases but organizational forms remain fixed.

Our empirical analysis finds that - in the absence of adverse weather and at average levels of airport congestion - majors using only integrated regionals at an airport experience average departure delays that are about 2.3 minutes shorter than those experienced by majors using only independent regionals. This performance advantage increases to about 5.6 minutes on days with "heavy" rain. It also increases with airport congestion. These results emerge in both ordinary least squares (OLS) and instrumental variables (IV) models and are robust across alternative measures of on-time performance. When we try to capture adverse weather using measures of snowfall, the results are somewhat weaker. OLS specifications find that the performance advantage of integrated majors decreases on days with "heavy" snow.³ IV specifications find a negative point estimate, indicating an increase in the performance advantage, but in most specifications it is imprecisely estimated. However, when we examine cancellation rates and very long delays, we find that - for these particular margins - the performance advantage of integrated majors does increase on days with snow. Overall, our results indicate that

³ As we explain below, we expect the OLS estimate to be biased upwards due to the endogeneity of integration decisions.

integrated majors have systematically better operational performance than non-integrated majors and this that performance difference grows larger on days and at times when real-time adaptation decisions are likely to be needed.

In addition to being one of the only papers to provide evidence on the effects of vertical integration on operational performance, we believe this paper is also one of the few empirical contributions that focus explicitly on the relationship between integration and adaptation decisions and the first that does so without relying on survey measures.⁴ Williamson (1975, 1985) first developed the hypothesis that integration facilitates adaptation decisions. Bajari and Tadelis (2001) and Tadelis (2002) further develop the idea that the need for *ex post* adaptation decisions can be a source of transaction costs and can therefore influence both contract design and integration decisions. In our previous work, Forbes and Lederman (2009), we investigate the determinants of vertical integration between major and regional airlines. Our focus there is also on the need for *ex post* adaptation and we show that majors are more likely to use owned regionals on routes on which they expect to have to make more adaptation decisions and on routes on which having adaptation decisions resolved sub-optimally is more costly. This paper builds on our earlier work by focusing explicitly on identifying the causal effect of integration on an important performance metric in this industry and providing evidence that sheds light on the underlying mechanism.

There is an existing set of papers that focuses on the consequences of vertical integration. Lafontaine and Slade (2007) provide a survey of this literature. Within this literature, one can distinguish between those papers that test market power based theories of integration (such as Hortacsu and Syverson, 2007) and those that – like us - test incomplete contracting based theories. The latter category includes papers such as Mullainathan and Scharfstein (2003) and Ciliberto (2006) both of which investigate the effects of organizational form on investment decisions as well as Baker and Hubbard (2004) and Gil (2007) which, respectively, explore the implications of organizational

⁴ There are empirical studies which test whether complexity (which makes complete contracts more difficult to write) or asset specificity (which makes the resolution of conflicts arising under incomplete contracts more costly) affect the likelihood that a transaction is organized internally. See, among others, Monteverde and Teece (1982), Anderson and Schmittlein (1984), Masten (1984), Masten and Crocker (1985), Joskow (1985), Hubbard (2001) and Levin and Tadelis (forthcoming). See Lafontaine and Slade (2007) for one review of this literature and Hubbard (2008) for a more detailed discussion of the evolution of this empirical literature.

form decisions in the trucking and movie industries. Our paper is perhaps most closely related to Novak and Stern (2008) which also examines the relationship between vertical integration and a specific performance measure – in their case, *Consumer Reports* ratings of automobile systems. They find that integrated firms have lower initial ratings, but greater improvements in ratings over time, which can be interpreted as evidence that integration facilitates adaptation when changes become necessary.⁵

This paper is organized as follows. In Section II, we explain why ownership of a regional may lead to operational efficiencies, particularly when non-contractible schedule adjustments are needed. Section III describes our empirical approach and Section IV addresses data and measurement issues. We present our results in Section V. In Section VI, we provide some additional discussion. A final section concludes.

II. Vertical Integration in the Airline Industry

II.A. The Role of Regional Airlines⁶

Regional airline service represents a large and growing fraction of U.S. domestic air travel. In our sample period, about one out of every seven domestic passengers was traveling on a regional carrier. Regionals operate as “subcontractors” for large network carriers on low-density short and medium-haul routes. These are routes which are most efficiently served with small aircraft.⁷ Majors subcontract these routes because regionals have a cost advantage in operating small aircraft resulting from the substantially lower compensation that regional airline employees receive.⁸ Majors do not typically operate any small aircraft themselves; thus, the decision whether to use a regional to serve a route is effectively a decision about plane size.

Regionals operating as subcontractors for majors do so under codeshare agreements. Under these agreements, the regional operates flights on behalf of the major.

⁵ There is also a set of papers which 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 mostly from the strategy field.

⁶ For a detailed description of the role of regionals in the U.S. airline industry, see Forbes and Lederman (2007).

⁷ Examples of such routes include Boston to Burlington, VT, or New York City to Albany, NY.

⁸ See Forbes and Lederman (2009) for a discussion of the source of lower labor costs among regional airline employees. In addition, Hirsch (2007) contains a comparison of pay rates at majors and regionals.

These flights are operated under the major's flight designator code and are marketed under the major's brand. For example, the regional Comair operates flights for Delta under the name Delta Connection. Many regional flights operating under codeshare agreements connect small airports to one of the major's hubs and many of the passengers travelling on the regional connect to or from flights operated by the major. The major is responsible for the scheduling, marketing and ticketing of the regional's flights, even those that do not connect to one of its hubs. Thus, regionals are providing an input – specifically, the operation of the aircraft - into a final product that is sold by the major. This is the sense in which majors and regionals are involved in a vertical relationship.

II.B. Organizational Forms

Relationships between majors and regionals are governed by one of two organizational forms – (1) a regional may be independently owned and contract with one or more major carriers; or (2) a regional may be wholly-owned by the major with which it partners.⁹ Table 1 lists the major-regional partnerships that were in place in 2000 for the carriers in our sample. Regional carriers that appear in bold were fully owned by their major partner. The table shows that there is substantial heterogeneity both across and within majors in the extent to which regional partners are owned.

In the case of an *owned regional*, the major carrier owns the assets of the regional but the regional and the major technically maintain separate operations. The main reason they separate their operations is so that they can maintain distinct labor contracts (one for the major's own employees and one for each of its regional's employees) and thereby preserve the cost advantages that regionals provide.¹⁰ We use the term “vertical integration” to refer to this relationship between a major and an owned regional.

The relationship between majors and *independent regionals* is governed by contracts. In general, these contracts specify which routes the regional will serve for the major, the planes that the regional will use and the schedule of flights. Contracts between majors and independent regionals take one of two forms. Historically, most were

⁹ In which case, we do not observe that regional operate flights for competitors of its parent company.

¹⁰ 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 the application is granted, the unions of the carriers will operate as a single entity.

revenue-sharing agreements under which the major and the regional shared the revenue from passengers whose itineraries involved travel on both airlines. Beginning in the late 1990s, majors and regionals have increasingly used what are known as “capacity purchase agreements”.¹¹ Under these agreements, the major retains all ticket revenue and pays the regional a fixed amount on a block-hour or flight-hour basis. The payment is structured to cover the regional’s costs and provide it with a reasonable rate of return. Capacity purchase agreements insulate a regional from revenue risk but leave it the residual claimant on profit increases that result from effective management of costs such as crew wages and lodging expenses. Relevant to our analysis here is the fact that, under both contract types, independent regionals face financial incentives that are based only on the routes that they serve and not on the remainder of the major’s network.

II.C. Vertical Integration and Operational Performance

We hypothesize that performance benefits to ownership might arise in this setting because contracts between majors and independent regionals are incomplete on at least one important dimension – real-time schedule adjustments. Schedule disruptions are common in the airline industry, resulting from a variety of factors such as adverse weather or mechanical problems. When they arise, majors may have to reschedule their own flights as well as the flights operated by their regional partners. Inclusion of these types of schedule adjustments in a contract is unlikely to be feasible because the contract would have to specify *ex ante* the full set of changes that would be made under every possible contingency and the precise manner in which the regional would carry out these changes.¹² Moreover, because the contracts used in this industry do not tie the regional’s compensation to the performance of the major’s network, an independent regional has no financial incentives to comply with schedule changes that may not be in its own

¹¹ This discussion draws on American Institute of Certified Public Accountants (2007) which provides more detail on the various contractual forms in this industry. See also Forbes and Lederman (2010) for a more detailed discussion of the change in contractual form and an analysis of the incentive properties of the two contract types.

¹² Even if it were feasible to specify these contracts, the fact that such contracts are not written suggests that it would be prohibitively costly to do so. Carlton (1979) shows theoretically that reducing uncertainty over the availability of inputs can be a motivation for vertically integration.

interest.¹³ Schedule changes that the major requests may even impose costs on the regional. For example, schedule changes may require a regional to pay overtime to its crew and independent regionals are responsible for their labor costs. Schedule changes ordered by the major may also negatively impact the regional's own performance statistics which could be important in seeking new business from other major carriers. For these reasons, independent regionals may have limited incentives to execute the real-time schedule changes that their majors request

Why would this impact the operational performance of the major? At a given airport, an airline's own flights and its regional flights are integrated into a common network and compete for access to potentially scarce airport and airline resources. Takeoff and landing slots, ground crew, equipment like de-icing machines and sometimes gates can be shared across majors and regionals. When schedule adjustments become necessary, a major trying to optimize its network may ask its regional to take certain actions that give the major's flights preferential access to these resources. Whether and how these actions are carried out by the regional can impact the performance of the major's flights. Consider the following examples: (1) When an airport is affected by adverse weather, airport activity will be reduced because more time is required between takeoffs and landings. Airlines will likely need to delay a large fraction of their flights. A major may choose to delay its regional flights so that larger, fuller mainline flights can depart and may ask a regional to move its aircraft from its gates. The willingness and speed with which the regional does this will impact the major's ability to avoid further delays on its own flights. (2) When a flight lands later than scheduled, the gate and/or ground crew scheduled for this flight may no longer be available. To avoid a lengthy arrival delay as well as a lengthy departure delay on the aircraft's next leg, a major may try to "turn the plane around" more quickly than usual. The major may ask the regional's employees to help unload and load baggage from the plane. The willingness and speed with which the regional does this will clearly affect delays experienced on this flight. More generally, if airport or airline resources are scarce, then any time flight departures

¹³ While capacity purchase agreements do separate a regional's incentive to operate a flight from the stand-alone profits of that flight, they still do not provide a regional with incentives to behave in a way that maximizes the overall profitability of the major's network. In Section VI, we address the question of why majors and regionals are unlikely to be able to write contracts that would address this incentive problem.

or arrivals deviate from their set schedule, integration may allow a major to more efficiently respond to the schedule disruption by improving its access to resources that had been allocated to the regional. It is worth emphasizing that while access to a certain set of shared resources may give rise to performance implications for the major's flights, aircraft and crew are not shared across majors and regionals. Thus, they are not the source of the externalities that we are trying to measure.

Why would employees at owned regionals have better incentives to comply with the major's requests? Owned regionals are likely to be less concerned with their financial and operational performance metrics. Furthermore, owned regionals should be more concerned about the profits and overall financial health of their major. Because they fly only for the major that owns them, if that major were in financial difficulty and had to divest itself of its regional unit, this could impose costs on the regional's employees who, after divestiture, might have to accept the lower salaries that prevail at independent regionals. In addition, career paths between owned regionals and majors are better established.

In Table 2, we investigate whether the raw data show any indication of a performance difference between majors using owned and majors using independent regionals. For the purposes of this simple table, we divide majors into two groups – those that have fewer than or exactly 50% of their regional flights at an airport operated by an owned regional and those that have more than 50% of their regional flights at an airport operated by an owned regional. We compare the mean departure delay, the 25th, 50th and 75th percentile departure delay and the cancellation rates of these two groups. The table shows that the mean departure delay on majors' flights departing from airports at which they use owned regionals is about 3 minutes shorter than the mean delay on flights operated by majors using independent regionals. Interestingly, there appears to be no difference in the 25th and 50th percentiles of their delay distributions but a four minute difference in the 75th percentile. Cancellation rates also appear to be lower for flights operated by majors using owned regionals. While this is, of course, very preliminary evidence, it does suggest that there may be systematic performance differences across integrated and non-integrated firms in this setting. We now turn to a more formal empirical analysis.

III. Empirical Approach

Ideally, we would like to investigate whether ownership of a regional allows a major to come closer to its profit-maximizing departure schedule, given the realized values of all relevant variables (for example, weather and air traffic control problems). Clearly we cannot calculate what an airline's ideal departure schedule would be in every state of the world. We can, however, measure a major's delays and cancellations on each flight that it operates. Therefore, for our empirical analysis, we assume that fewer delays on average (relative to the major's original departure schedule) indicate that a major is coming closer to achieving its optimal departure schedule. We also explore whether there is evidence of differences in very long delays and cancellation rates. While we would also like to explore the impact of ownership on the regional's performance, data limitations prevent us from doing so because most regionals are too small to meet the reporting requirements of the Bureau of Transportation Statistics. Our results with respect to the major's own flights should therefore be considered only a partial estimate of the effects of integration in this setting. However, given the substantially larger size of majors compared to regionals (majors carried six times as many passengers as regionals in 2000), our estimates should reflect a large portion of the total effect.

III.A. Empirical Specification

Our empirical specification regresses a major's performance on a particular flight on its extent of integration with the regional carrier(s) that it uses at the origin airport of that flight. We exploit the fact that there are many airports at which some majors use owned regionals while others use independent regionals and include fixed effects for each origin airport-day combination in our model. Thus, we are able to test whether – *at a given airport, on a given day* - the operational performance of majors using owned regionals differs from that of majors using independent regionals.

As described in the previous section, we hypothesize that ownership may mitigate an incentive problem that is particularly likely to arise when majors need to make non-contractible changes to a regional's schedule. To investigate whether the performance effects of ownership may result from the fact that ownership of a regional facilitates real-

time schedule adjustments, we identify situations in which a major is more likely to have to make unanticipated changes to its set schedule of flights. We do this in two ways. First, we exploit the fact that adverse weather is one of the leading causes of schedule changes and use measures of the daily weather at an airport as proxies for the likelihood that flights departing from that airport will be affected by non-contracted schedule adjustments. Second, we construct a measure for airport capacity utilization to identify times of day during which airport congestion may cause an airline to have to make schedule changes.¹⁴ We interact our integration measure with measures of daily weather and hourly congestion to test whether ownership has a different effect on operational performance on days with particularly adverse weather conditions or at congested times of day.

Specifically, we estimate the following equation:

$$PERF_{fir}^t = \alpha_o^t + \delta_1 OWNED_{ir} + \delta_2 OWNED_{ir} * ADAPT_{fr}^t + X_{fir}^t \beta + \varepsilon_{ifr}^t \quad (1)$$

where $PERF_{fir}^t$ is a measure of airline i 's operational performance on flight f on route r on day t , α_o^t is an origin airport-date fixed effect, $OWNED_{ir}$ measures the extent of airline i 's ownership of its regionals serving the origin airport of route r , $ADAPT_{fr}^t$ is a vector of variables that measure the extent of adverse weather at the origin airport of route r on day t or the level of congestion at the origin airport of route r around flight f 's scheduled departure, X_{fir}^t is a vector of control variables (including flight, airline and airport level variables), and ε_{ifr}^t is an error term. If there are performance benefits to ownership and if these are greater on days with adverse weather or during more congested times of day, then we would find both $\delta_1 < 0$ and $\delta_2 < 0$.

¹⁴ We thank an anonymous referee for this suggestion. While airlines cannot anticipate the actual weather on a given day at the time that they schedule the flight, expected airport capacity utilization is known *ex ante*. One could therefore worry that the major's decision whether to schedule a flight during a congested time period is endogenous to the performance equation. However, most of the flights in our data involve a hub endpoint and thus scheduling decisions of both the hub and non-hub carriers at those airports will be heavily influenced by hub banking times as well as the underlying distribution of demand over the day. The mean value of our congestion variable is quite similar across integrated and non-integrated regionals suggesting they do not systemically schedule at more or less congested times. A previous version of the paper estimates our model without the congestion variable and finds similar results on the direct effect of vertical integration and its interaction with weather.

III.B. Endogeneity

Because ownership decisions are made by optimizing firms, measures of integration will typically be endogenous in a performance equation (see Masten, 1993, and Gibbons, 2005). A valid instrument must be correlated with the firm's ownership decision but uncorrelated with the error term in the performance equation. However, because the same variables determining the ownership decision will also appear in the performance equation, instruments for organizational form decisions are difficult to find. While this is true in our setting as well, the fact that we are measuring the impact of ownership on a *specific* performance margin rather than the overall performance of the firm gives rise to a potential source of instruments which will valid under certain conditions.

Based on Forbes and Lederman (2009), one can think of a major's decision whether to use an owned regional on a particular route as depending on: (1) the returns to integration on the regional route; (2) the returns to integration on other routes in the major's network (especially those that share an endpoint with the regional route); and (3) the returns to integration that accrue systemwide (for example, through higher labor costs). Our performance equation measures whether a major's ownership of the regional(s) that it uses at a particular airport affect the performance of its flights that depart from that airport. Variables that influence the returns to integration elsewhere in the major's network - for example, on the regional routes themselves - but do not directly influence the major's performance on the flight in question will be valid instruments for the major's integration decision at the origin airport of that flight.

The logic of this instrumental variables approach is best illustrated with an example (see Figure 1 for a representation of this example). Consider a particular Delta Air Lines flight from Boston to Atlanta one day in 2000. This represents a single observation in our data. For this observation, our ownership variable measures the extent to which Delta uses an owned regional to serve its regional routes into and out of the Boston airport, routes such as Boston-Albany or Boston-Burlington. Delta's decision whether to use an owned or independent regional on, say, the Boston-Burlington route will depend on the returns to ownership on the Boston-Burlington route, on Delta's other routes out of Burlington and on Delta's other routes out of Boston, with the latter being

precisely what our performance equation is trying to measure. Characteristics of the Burlington airport will affect the returns to integration on the Boston-Burlington route as well as other routes out of Burlington and, as we show in our earlier work, are strongly correlated with Delta's ownership decision on this regional route. Assuming characteristics of the Burlington airport do not directly affect Delta's performance on the Boston-Atlanta route, they can serve as valid instruments for the ownership variable. The same is true for the characteristics of the other regional endpoints that Delta connects to Boston with a regional partner. Based on this logic, we instrument for major's ownership decision at a particular airport with the characteristics of the endpoint airports that its regional(s) connect to that airport. In particular, we use the characteristics that our earlier work found to predict owned regional use. These characteristics are the long-run average precipitation and snowfall at the airport and whether the airport is a hub for the major.¹⁵ Even though the returns to integration might vary with short-run fluctuations in weather, our earlier work and hence our instruments here use the long-run average weather at an airport because we assume adjustment costs prevent firms from changing their integration status in the short-run.

Continuing with our example, the validity of our instruments depends on the assumption that the long-run weather conditions at Burlington and its hub status are uncorrelated with the error term in the performance equation for Delta's Boston-Atlanta flight. In order to assess the reasonableness of this assumption, it is helpful to think through what might be contained in the error term of our performance equation. Even though our models include a full set of departure airport-date fixed effects and a large number of control variables, we only explain about 16 percent of the variation in departure delays. The remaining variation must be due to airport-specific factors that vary *within* the day (and so are not captured by our airport-date fixed effects), airline-airport specific factors that vary over the course of the year and idiosyncratic flight-level unobservables. Such factors can include the actual timing of bad weather, air traffic control problems that occur at a specific time of day, mechanical or crew problems that

¹⁵ The precise construction of the instruments and the data used for the instruments are described in Appendix A. One might believe that it is the variance of weather conditions, not their mean, which predicts integration. We have explored this and found that the mean and the variance are highly correlated and we cannot separately identify their effects.

affect a specific flight, or problems that occur elsewhere in the airline's network that lead to propagating delays.

Because our instruments measure characteristics of other airports in the major's network, the risk of propagating delays could result in correlation between our instruments and the error in the performance equation. For example, one might be concerned that long-run weather conditions at the Burlington airport are correlated with realized weather conditions at Burlington and that realized weather at Burlington on a given day is correlated with Delta's delays on flights from Boston that day. For a number of reasons, we believe that long-run average weather conditions at the regional endpoints are unlikely to be correlated with the error terms in the performance equation. First, we estimate specifications in which we add explicit controls for the realized weather at the regional's endpoint airports on the same day and find that doing so has no effect on our two-stage least squared estimates. Second, the correlation between average historical weather and the amount of rain or snow on a given day is very low.¹⁶ Third, while delays certainly propagate within the major's network, delays are much less likely to propagate from the regional's network to the major's network because aircraft and crew are not shared across majors and regionals (i.e. the primary mechanism through which delays cascade from one flight to another does not operate here). Although passengers may connect between regional and major flights, airlines typically do not "hold" outgoing flights for late incoming passengers unless it is the last flight of the day which is something we can directly control for and find that doing so does not change the results.

One could also raise a concern about correlation between the hub instrument and the error term in the performance equation. Specifically, one could argue that whether or not Burlington is a hub for Delta provides information about the role of Boston in Delta's network. For example, regionals do not typically fly interhub routes. So, if Burlington were a hub for Delta that would indicate that Boston was not. If we could not perfectly control for the relationship between an airport's role in an airline's network and its delays, this could lead to correlation between the hub instrument and the error term in the performance equation. However, as we explain in Section IV, we include a number of airline-airport level control variables in the model, among them controls for whether

¹⁶ We list these correlations in Appendix A.

either endpoint of the route is a hub for the major. As well, we also estimate specifications with airline-airport fixed effects. These will explicitly capture any relationship between the major's network configuration and its departure delays. While these specifications do not allow us to estimate the direct effect of integration on performance, we can still estimate the interaction terms with weather and congestion, and we find that the results on these terms are highly robust to including the airline-airport fixed effects.

Finally, to the extent that there is any residual correlation between average characteristics of the endpoints served by a regional and the performance of a major on its own flights on a given day, it should bias our two-stage least squares results towards finding that majors using owned regionals perform *worse* since our previous work shows that that airlines vertically integrate on routes that are *more* likely to experience schedule disruptions.

IV. Data and Measurement

IV.A. Data Sources

Our primary source of data 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.¹⁷ We augment these data with information from several other sources. Data from the Official Airline Guide (OAG) provide the complete flight schedules of all domestic airlines, regionals as well as majors.¹⁸ Data from the Regional Airline Association (RAA) indicates 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. Data on the daily weather at each airport are taken from the National Oceanographic and Atmospheric Administration (NOAA).

IV.B. Construction of the Sample

¹⁷ Carriers are required to report these data if they account for at least one percent of domestic passenger revenues in the prior year.

¹⁸ Our data provide a representative week for each quarter.

Our sample includes domestic flights operated by the seven largest network carriers (American, Continental, Delta, Northwest, TWA, United and US Airways) in the year 2000.^{19,20} We begin with their flights that depart from the largest 100 U.S. airports and then impose the following restrictions.²¹ First, we exclude flights that depart from or arrive at airports in Alaska, Hawaii, Puerto Rico, Guam or the U.S. Virgin Islands because the nature of routes to these states and territories is quite different from travel within the contiguous 48 states. Second, because our empirical approach exploits variation *across* airlines at an airport, we exclude departure airports at which we do not observe at least two majors using a regional. Third, 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, 2008, for details). Fourth, we drop observations with missing daily weather data. 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 a major’s flights from airports at which it does not use a regional at all. Finally, we exclude flights on Saturdays and Sundays so that our within-airport variation in an airline’s extent of vertical integration is not driven by within-week fluctuations in regional use on the same route. Our final dataset includes 1,981,807 flights departing from 72 departure airports and arriving at 159 arrival airports on 260 days.

IV.C. Variables

Variable names and definitions appear in Table 3. Summary statistics are in Table 4.

i. Performance Measures

¹⁹ TWA went through financial difficulties during this period and was acquired by American Airlines in April, 2001. We have checked that our results are robust to excluding TWA.

²⁰ All of the traditional network carriers employ regionals to some extent. The so-called “low-cost carriers”, such as Southwest Airlines, do not subcontract flights to regional carriers.

²¹ Airport rankings are based on year 2000 enplanements and are compiled by the Federal Aviation Administration. Airport size decreases quite rapidly. For example, the largest airport based on enplanements is Atlanta (39.3 million passengers), the 20th largest airport is Philadelphia (12.3 million passengers), the 50th largest is San Antonio (3.5 million) and the 100th largest is Harrisburg (644,000).

Our main dependent variable is *Departure Delay* which measures the time between the scheduled departure and the actual departure of an aircraft from the gate. This is our preferred measure of performance because – compared to delays incurred on the runway or in the air - delays incurred at the gate are more likely to be under the airline’s control. Note that the departure delay will be negative if the aircraft pushes back early from the gate. Short negative delays of a few minutes are quite common. However, we exclude flights that leave more than 15 minutes early, as we suspect that they may represent a rescheduled flight. We also exclude very long delays that appear to be flights that have been delayed and rescheduled for departure on the following day. To do this, we identify flights for which the actual departure time recorded is *earlier* than the scheduled departure time and the minutes of delay recorded are consistent with the flight departing at the earlier departure time on the subsequent day.²² We exclude these flights because part of the delay incurred in these cases would be attributable to the fact that airports do not operate overnight, as opposed to the airline’s poor performance.²³ As reported in Table 4, the average departure delay in our sample is just over 11 minutes.

The empirical analysis also uses several alternative dependent variables. *Cancelled* is a dummy variable that equals one if the flight is cancelled. *Arrival Delay* measures the difference between a flight’s actual arrival time and its scheduled arrival time. *Cancelled or Arrival Delay > 30 minutes* is a dummy variable that equals one if a flight is more than 30 minutes late on arrival or cancelled. We construct an analogous variable using a 60 minute arrival delay. The summary statistics in Table 4 indicate that about 4% of flights in our sample are cancelled, 18% are delayed over 30 minutes or cancelled and 11% are delayed over 60 minutes or cancelled.

ii. Ownership Measures

To measure the extent of a major’s vertical integration with its regionals at an airport, we measure the fraction of all regional flights that a major has departing from an

²² This drops 495 observations.

²³ We do not drop flights that depart on the subsequent day if they depart during the few hours after midnight during which most airports are still open.

airport on a day that are operated by a regional that is owned.²⁴ We call this variable *Fraction Owned*. As Table 4 indicates, the mean of *Fraction Owned* is 0.56. During our sample period, there are no changes in the ownership of any regional carriers. Almost all of the variation in *Fraction Owned* therefore comes from differences across the 282 carrier-airport combinations in our data.²⁵ The small amount of variation in *Fraction Owned* within carrier-airport over time arises from two sources. First, there are some changes in the number of regional flights by majors who use two regionals of different ownership types at the same airport. Second, there are four instances in which a carrier switches from using one type of regional at an airport to using a different type of regional at that airport. These are all cases in which the major had an existing relationship with the new regional at other airports prior to the switch. Our results are robust to excluding these carrier-airport combinations from the estimation.

As an alternate measure of integration, we also construct *Owned Regional Flights* which measures the number - rather than fraction - of all regional flights that a major has departing from an airport on a day that are operated by a regional that is owned. While this measure has the advantage of capturing the scale of the regional's operations, a drawback of this measure is that our instruments are only able to predict the decision to vertically integrate, not the scale of regional operations. As a result, we only use this variable in OLS specifications. In our sample, the average number of owned regional flights that a major has operating at an airport on a day is 65.

iii. Proxies for the Likelihood of Adaptation (Weather and Congestion Measures)

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 *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

²⁴ 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.

²⁵ An analysis of variance shows that these differences explain 99% of the variation in *Fraction Owned*.

degrees Fahrenheit or less.²⁶ The average daily rainfall in our sample is 0.11 inches and the average daily snowfall is 0.08 inches. Of course, there are many days on which there is no snow and many airports for which there is never any snow.

Our empirical approach requires us to measure “adverse” weather – i.e.: weather conditions that are likely to necessitate schedule adjustments. To construct our main measure of adverse weather, we first calculate the 95th percentile of the daily rain distribution for each airport in our sample. We then construct the dummy variable *Rain>95th Percentile* which equals one on days on which the observed rainfall at the airport exceeds the 95th percentile of that airport’s rain distribution. Thus, roughly speaking, *Rain>95th Percentile* captures an airport’s 18 rainiest days of the year. We construct *Snow>95th Percentile* analogously. The average amount of rain on days with *Rain>95th Percentile* equal to one is 1.29 inches and the average amount of snow on days with *Snow>95th Percentile* equal to one is 3.11 inches. We construct all of the weather variables for the both the departure and arrival airport of a flight.

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 ensures that – at least for the rain variable - bad weather events are observed at all airports in our sample. Offsetting these benefits is the drawback that these measures do not capture the severity of the weather in any absolute sense. Therefore, we also present specifications where we use the linear *Rain* and *Snow* variables and specifications where we use dummy variables that indicate the presence of any rain (*Rain>0*) or any snow (*Snow>0*).

In addition to adverse weather, the likelihood of needing to make adaptation decisions may be greater during more congested time periods. To capture this, we construct a measure of airport congestion that varies over the course of a day. We do this by first calculating the maximum number of flights that we observe arriving at or

²⁶ 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.

departing from an airport in an hour at any point in our sample. We take this as a measure of the airport's capacity. Then, we calculate the number of flights arriving at or departing from an airport during any given hour and divide that by our capacity measure. The resulting ratio is called *Congestion* and it is matched to the flight-level delay data at the level of the departure airport-hour.²⁷ The mean of this variable in our sample is 0.66. Note that in our regressions, we measure this variable relative to its sample mean so that the coefficient on the uninteracted *Fraction Owned* variable can be interpreted as the effect of integration at the average level of congestion as opposed to at zero congestion.

iv. Flight Level Control Variables

Because delays are generally thought to worsen over the course of a day, we control for the departure time of a flight. To do this, we divide the day into seven three-hour blocks and construct a dummy variable for each of these blocks. We use the 5-8am block as the excluded time block.

v. Airline-Airport Control Variables

We control for whether a major's flight departs from one of its hub airports (*Departs from Hub*) or arrives at one of its hub airports (*Arrives at Hub*). We construct *Regional Flights* which equals the total number of daily regional flights that a major has departing from an airport and use this to control for an airline's scale of regional operations (of either type) at an airport. We also interact these controls with the weather measures to ensure that our interactions between the ownership and weather measures are not capturing the fact that adverse weather may differentially affect airlines with a larger scale of operations (or a hub) at a particular airport.

vi. Airport Control Variables

We construct *Airport Flights* which measures the total number of domestic flights scheduled to depart from (arrive at) an airport on a given day. We also construct *Slot*

²⁷ For 31 of the airports in our sample, we have estimates of the airport's capacity from the Federal Aviation Administration. We have calculated an alternative congestion variable using this as the denominator. For these 31 airports, the two congestion measures have a correlation of 0.82 and the regression results are very similar. However, we prefer to use the one described in the text as it allows us to measure congestion for our full sample.

which is a dummy for whether the airport is slot-controlled.²⁸ We construct these for both the departure and arrival airports of a flight. However, in most specifications, the departure airport variables will not be separately identified from the fixed effects. 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

V. Results

Our presentation of the results proceeds in several stages. We begin by briefly describing the first stage regression for *Fraction Owned*. We then describe a series of OLS regressions that are presented in Table 5. These regressions present the basic relationships in the data, gradually incorporating additional fixed effects. After that, we discuss the results of estimating our model using the instrumental variables approach described above. We then present several extensions using alternate performance metrics and weather measures. We conclude with some back-of-the-envelope calculations.

V.A. First Stage Regression

Appendix B presents the results of the first stage regression of *Fraction Owned* on the excluded instruments and the exogenous variables. All of the instruments have highly significant effects and the signs of the effects are consistent with the findings in Forbes and Lederman (2009). Owned regionals are more likely to be used when a greater fraction of the regional's routes connect to the major's hubs and when the endpoints served by the regional experience greater annual rain and snowfall. Endpoints with more months with below freezing temperatures are less likely to be served by owned regionals.²⁹ Joint significance of the instruments is confirmed by the F-statistic presented at the bottom of the table. The overall R-squared of the regression is 0.58 while the within R-squared is 0.13.

²⁸ 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).

²⁹ In the earlier paper, we explain that this result is consistent with the observation that those airports have shorter delays on average.

V.B. Ordinary Least Squares Results

Table 5 presents a series of OLS regressions. In the first column, we regress *Departure Delay* on *Fraction Owned* and our control variables. We do not include fixed effects in this specification so that we can show the coefficients on all of the control variables. Consistent with the raw data in Table 2, we find that vertical integration between a major and regional is associated with shorter delays on the major's flights. The point estimate implies that a major using only integrated regionals at an airport has departure delays that are, on average, three minutes shorter than a major using only independent regionals. This compares to a mean delay in the sample of about 11 minutes. Thus, this initial finding is both statistically and economically significant.

In the second column of the table, we add fixed effects for each airport-date combination and include these fixed effects in all subsequent specifications. The estimate on *Fraction Owned* is virtually unchanged. The third column of the table adds interactions of *Fraction Owned* with the variables that proxy for the likelihood of having to make adaptation decisions (*Rain*>95th *Percentile*, *Snow*>95th *Percentile* and *Congestion*). Consistent with the discussion in above, the results indicate that the performance advantage of vertically integrated majors increases on days with "heavy" rain and during congested periods of the day. However, we find that the performance advantage actually decreases on days with "heavy" snow. The magnitudes of the estimates suggest that – on days with "good" weather and during times with average levels of congestion – majors using only owned regionals at an airport have departure delays that are about 3 minutes shorter than majors using only independent regionals. On days with rain above the airport's 95th percentile, this advantage increases to about 7 minutes. As airport congestion increases, the performance advantage also increases, by about 0.4 of a minute for every 10 percentage point increase in congestion. The coefficient on the snow interaction suggests that the performance advantage of integrated majors is entirely eliminated on days with snowfall above the airport's 95th percentile.

The next two columns of the table add additional fixed effects. In (5-4), we add airline fixed effects to control for systematic differences in delays across the seven airlines in our sample. Once we include these fixed effects, the coefficient on *Fraction Owned* becomes harder to identify because only four of our seven airlines use both

owned and independent regionals. We do not estimate a significant coefficient on *Fraction Owned* in this specification but all of the interaction terms have the same signs as before and are still significant. In (5-5), we replace the airline fixed effects with airline-airport fixed effects. These fixed effects completely absorb the uninteracted ownership measure but still allow the interaction terms to be identified. The results on these interaction terms are again consistent with those in the previous columns.

In the final column of Table 5, we re-estimate (5-3) using *Owned Regional Flights*. Since we already include the number of regional flights of any type as a control, this variable measures the impact of increasing the number of flights by an owned regional carrier by one, while holding the major's overall number of regional flights constant. This variable has the advantage that it allows for the scale of integrated regional operations to matter. We find that the OLS results are robust to this alternative way of measuring integration. Carriers with a greater number of flights operated by an owned regional have a performance advantage which increases on days with heavy rain and at congested times of the day. The interaction with extreme snow is positive, as we find in all previous columns. A comparison of the estimates on *Fraction Owned* and *Owned Regional Flights* in (5-3) and (5-6) shows that the magnitudes of the estimated effects are quite similar. The average number of owned regional flights that a major has an airport at which it uses an owned regional at all is about 100. The coefficient on *Owned Regional Flights* (which is measured in hundreds) is -3.55 which indicates that, at the mean, the performance advantage is about three and a half minutes. This is very similar to the coefficient of -2.935 estimated on *Fraction Owned* in (5-3).

With respect to the coefficients on the control variables, we find that delays are longer for flights departing during congested times and for flights that depart later in the day. We find that delays are generally longer for flights that depart from or arrive at the airline's own hub and for flights that depart from or arrive at larger airports or slot-controlled airports. Delays are also longer on days with rain or snow above the 95th percentile of the airport's distribution (with weather conditions at both the departure and arrival airport having an effect a flight's departure delay). We also find that coefficients on the interaction effects of the airport size variables and the weather variables are positive indicating that adverse weather leads to even longer delays at busier airports. The

effects of most of the control variables are consistent across specifications although, not surprisingly, some decrease in magnitude or even change signs once we include the airport-date fixed effects or the additional airline or airline-airport fixed effects.

Our results on the control variables are consistent with previous work on the determinants of flight delays, such as Mayer and Sinai (2003), Mazzeo (2003), Rupp *et al.* (2006), Ater (2009) and Rupp (2009). This work has largely focused on two empirical questions: (1) the relationship between hubs and delays; and (2) the relationship between route-level competition and delays. In perhaps the most well known empirical paper on delays, Mayer and Sinai find that hub airports have longer delays than non-hub airports and that, at hub airports, the hub carrier has disproportionately longer delays than non-hub carriers. This latter finding is consistent with our finding of a positive coefficient on the *Departs from Hub* variable in regressions that include airport-date fixed effects. Rupp *et al.* and Rupp also find longer delays for flights to or from hubs while Mazzeo finds that delays for such flights are shorter. With respect to congestion, like us, Mazzeo and Ater find that delays increase with congestion. Like us, Mazzeo also finds that delays increase with the time of day. Finally, with respect to weather, Mazzeo, Rupp *et al.* and Ater include weather controls in their regressions and find, as expected, that delays are longer in worse weather conditions.³⁰

V.C Instrumental Variables Results

Table 6 presents the results using the instrumental variables strategy described above. All specifications include airport-date fixed effects and all of the control variables from Table 5; however, for the sake of space, the coefficients on some of these variables are not reported. The first column includes *Fraction Owned* but not its interactions with the weather and congestion variables. The IV results in (6-1) indicate a 2.8 minute performance advantage to integrated majors, virtually identical to the OLS finding in (5-2). Using the results from Forbes (2008) on the effect of flight delays on ticket prices, we

³⁰ Besides rain and snow, Mazzeo and Ater also include dummies for particular weather events, such as thunderstorms or fog. We have investigated these additional weather data, but found a very large number of missing values for the airports and the time period in our sample. In addition, the missing values did not appear to be random. We therefore decided not to include these variables in any way.

calculate that these longer delays would translate into an average reduction in the price of a roundtrip ticket of \$4.00 for direct passengers and \$2.17 for connecting passengers.³¹

The second column adds the interaction terms. The coefficient on the direct effect of *Fraction Owned* in this specification is again negative and implies that - on days with rain and snow below the 95th percentile of the airport's weather distribution and at the mean level of congestion - carriers that are fully integrated with their regionals at an airport have delays that are 2.4 minutes shorter than those that are not integrated. This performance advantage increases to about 5.6 minutes on days with rain above the 95th percentile of the airport's distribution. In addition, an increase in congestion of ten percentage points increases the performance advantage of integrated carriers by about one minute. It is interesting to note that while the estimates on *Fraction Owned* and its interaction with the rain measure are very similar to the OLS estimates, the coefficient on the congestion interaction doubles in magnitude. This suggests an upward bias of the OLS estimate which is consistent with the expectation that vertical integration is (optimally) used to govern more difficult transactions.

In contrast to the OLS results, in the IV specifications, the coefficients on the snow interactions are never significant and the point estimates are negative though very small in magnitude. This is again consistent with an upward bias in the OLS. While it may seem surprising that we do not estimate an additional performance advantage on snowy days, estimating the relationship with snow is difficult for several reasons. First, we have a large number of airports that never experience any snow and we have many airports that experience only small amounts of snow so that the 95th percentile of their snow distribution is either zero or a very small number. Second, it may be difficult to detect differences in how well different airlines at an airport deal with extreme snow conditions because these may simply shut down airports for periods of time. Finally, it may be the case that snow has more of an effect on long delays and/or cancellations. We explore these issues further when we investigate alternate measures of delay and alternate weather measures in Sections V.D and V.E below.

³¹ Note that Forbes' results are for arrival delays. To the extent that the equivalent arrival delays are longer than the departure delays - as our results below suggest - we would predict a larger effect on prices.

The remainder of Table 6 explores our IV results further. In column (6-3), we add controls for *realized* weather at other airports in the carrier's network on the same day. For example, for a major's flight between Boston and Atlanta on a particular day, we calculate the mean of *Rain*>95th *Percentile* and *Snow*>95th *Percentile* across the other airports that the major serves from Boston that day. We do the same thing for the endpoints served by the major's regional(s) from Boston that day. Controlling directly for realized weather at the endpoints that the major's regionals serve allows us to address the concern that the historical weather patterns at these airports – which we use as instruments – could be correlated with a source of error in the performance equation. The results in (6-3) show that the coefficients on the ownership variables are virtually unchanged when we add these controls. Furthermore, realized weather at the regional's endpoints has no effect on the delays experienced on the major's flights. In contrast, realized weather at the major's other endpoints has positive and relatively large effects. Consistent with the discussion in Section III, this suggests that delays are likely to propagate through the major's own network but less likely to propagate from regional's network to the major's because the major and regional do not share aircraft or crew.

In the final column of the table, we include airline-airport fixed effects. As described above, these additional fixed effects address the concern that there may be some residual correlation between our instruments and unobservable airline-airport characteristics. The results are again very consistent with what we have found earlier. Overall, the results in Table 6 clearly suggest that use of an owned - rather than independent – regional at an airport improves a major's operational performance at that airport, and that this performance advantage increases with extreme rain and with congestion.³²

V.D Alternative Performance Measures

Table 7 investigates the effects integration on four alternative measures of operational performance: *Cancelled*, *Arrival Delay*, *Arrival Delay*>30 or *Cancelled* and *Arrival Delay*>60 or *Cancelled* (see Table 3 for definitions). We use these additional performance measures for several reasons. First, we want to make sure that our baseline

results are robust to different ways of measuring flight delays. Second, we want to check that the shorter delays that majors using owned regionals experience on average are not coming at the cost of higher cancellation rates or an increase in very long delays. Third, we want to investigate whether integrated airlines have an advantage in terms of arrival delays, especially long arrival delays, which is presumably what consumers care most about.

In all four specifications in Table 7, we find that the direct effect of *Fraction Owned* is negative and highly significant and the coefficient estimates imply economically meaningful effects. From (7-1), we estimate that, on days without heavy rain or snow and at the mean level of congestion, integrated majors experience cancellation rates that are 1.7 percentage points lower than those of non-integrated carriers (the mean cancellation rate in our sample is about 4 percent). (7-2) indicates that integrated majors have arrival delays that are about 5 minutes shorter than non-integrated ones. (7-3) and (7-4) indicate that their likelihood of having a flight cancelled or delayed by more than 30 (60) minutes is 4.2 (2.1) percentage points lower.

The coefficients on the rain interactions are negative throughout the table, but are imprecisely estimated in (7-1) and (7-2). We also find negative coefficients on the snow interactions in all specifications except (7-2), and they are statistically significant in (7-1) and (7-4). These results suggest that the way in which integrated majors can affect their performance on days with extreme snow may be to reduce cancellations and long delays, even though the results in the previous table and in (7-2) show no statistically significant effects on mean departure and arrival delays. The congestion interactions are negative and highly significant throughout this table and their effects are large.

Besides providing a valuable robustness check, specifications (7-3) and (7-4) also allow us to explore whether our results might be driven by majors trying to “hide” delays by shifting them to their regionals, most of which are too small to meet government reporting requirements. While some shifting of delays to regionals is probably optimal (because regionals carry fewer passengers per departure and/or because they have lower labor costs), if all integration did was allow majors to shift delays to unobservable margins, it would not reflect a true efficiency. To investigate this, we exploit the fact that the statistic typically reported by the Department of Transportation and the media is the

proportion of a carrier's flights that are delayed more than 15 minutes. Therefore, if integrated majors were better able to shift delays to their regional partners, we expect that they would focus on flights with delays around 15 minutes and try to reduce these delays to just less than 15 minutes. We expect they would focus to a much lesser extent on flights with much longer delays. Our findings in Table 7 that integration also affects long delays and cancellations suggest that the benefits of integration are not merely due to the majors' efforts to manipulate on-time statistics.

V.E Alternative Weather Measures

Our final table explores alternative ways of measuring adverse weather, especially snowfall. In Table 8, we present one specification that replaces our previous weather measures with linear measures *Rain* and *Snow* and another one with dummy variables that equal one if the airport experienced any rain or snow on that day.³³ Our dependent variable in this table is *Arrival Delay > 60 or Cancelled* which is one of the performance measures that, in Table 7, revealed a potential performance advantage of integration on days with snow. We find again that the negative direct effect of *Fraction Owned* is very robust. For these specifications, we also find negative effects on all the interactions of *Fraction Owned* with rain, snow and congestion, although the rain effect is insignificant in (8-2). The negative coefficients on the snow interactions suggest that – at least for some performance metrics – the advantage of integrated regionals may increase on days with snowfall. However, we interpret even these results cautiously because these alternative snow measures still have insignificant effects on most of the other performance measures that we have used.

V.F Back of the Envelope Calculations

We conclude this section with simple back-of-the-envelope calculations that relate the performance advantage of integration that we find to a rough estimate of the costs of integration. We find that integrated carriers have departure delays that are on average 2.8 minutes shorter than those of independent carriers and arrival delays that are

³³ We have also estimated a robustness check where we redefine *Rain > 95th Percentile* and *Snow > 95th Percentile* to equal zero on days with very small amounts of rain or snow and the results are very similar.

on average 5.2 minutes shorter. Longer delays impose costs on airlines and reduce demand for their flights. While we do not have an estimate of how delays translate into higher costs, we can use Forbes' results on the relationship between arrival delays and ticket prices, and combine them with data on direct and connecting passengers on the routes in our sample³⁴ to estimate the demand-side benefit from integration for the average regional in our sample. We find this benefit to be about \$28 million a year. Since Forbes' results are for passengers traveling to and from LaGuardia Airport in New York City, who may have a higher value of time than the average U.S. airline traveler, this number is likely an upper bound. If we instead use the “[value] for aviation passenger travel time” of \$0.55 per minute (in 2000 dollars) that Federal Aviation Administration (1997) recommends based on a survey of travelers nationwide, we find a lower benefit of \$15 million a year.³⁵

How does this compare to the costs of integration? As Forbes and Lederman (2009) explain in detail, the most important component of these is the higher labor costs that result from vertical integration. While we do not have comprehensive data on differences in labor costs across regionals, we were able to get aggregate payroll data on five of the regionals in our sample for the year 2000 from the Bureau of Transportation Statistics. The total annual payroll for these regionals ranges from \$56 million to \$326 million, with an average value of \$157 million. We can only speculate on the actual differences in salaries between owned and independent regionals, but given these numbers those differences could certainly be of a similar order of magnitude as the benefits to integration that we compute.

VI. Discussion

Consistent with our discussion in Section II, our empirical results indicate that ownership may allow majors to more efficiently re-optimize their networks when schedule adjustments become necessary and airport or airline resources are scarce. While it is clear why the contracts used in this industry could lead to incentive problems between majors and independent regionals, one might question why either more

³⁴ These data come from the Department of Transportation's Databank 1B.

³⁵ Morrison and Whinston (1989) find that an increase of one percentage point in the share of flights delayed more than 15 minutes reduces passengers' willingness-to-pay by \$0.61, measured in 1983 dollars. Their estimates are not directly comparable to the delay measures we use here.

sophisticated contracts or informal contracts cannot be used to better align incentives. We believe that neither formal nor informal contracts are likely to resolve the incentive problem in this setting. First, consider a contract that compensates the regional based on the performance of the major's entire network. While this more closely aligns the regional's incentives with the major, the performance of the major's overall network depends on the major's effort as well as on the regional's effort and both efforts are likely unobservable to the other party. As a result, this contract would give rise to moral hazard problems and be unable to achieve the first-best outcome (see Holmström, 1982). Furthermore, such a contract would expose the regional carrier to a great deal of risk. Recent changes in the industry – in particular, a greater reliance on the capacity purchase agreements described in Section II - suggest that regional carriers should be regarded as risk-averse.

Second, consider a contractual arrangement that simply allocates to the major the rights to make any *ex post* adjustments to the regional's schedule. Our understanding is that this is occurs, at least under capacity purchase agreements. However, having the rights to order specific schedule changes is not equivalent to having the rights to implement those schedule changes.³⁶ At the time that schedule changes need to be executed, a major cannot simply replace its regional's labor and execute these changes itself. Since schedule changes ordered by the major must still be carried out by the regional, we expect that even with a contractual allocation of decision rights, incentive problems will remain.

Finally, consider the role that relational contracts may play. Within its relationship with a particular major, an independent regional may have an incentive to develop a reputation for cooperation if there is value in continuing the relationship with the major. However, cooperating with the major may lead to worse performance metrics for the regional. While the major it is working for will know whether the poor metrics reflect cooperation or poor performance, other majors that the regional may try to work for will likely not be able to distinguish the two. Cooperation may therefore have the effect of reducing the regional's ability to attract outside business which, in turn, will

³⁶ This is in contrast to the standard Grossman-Hart depiction of control rights which confer the right to determine precisely how an asset is used.

weaken its bargaining position vis-à-vis its current major. More generally, reputations for cooperation in this setting may be difficult to establish because the same observable metrics that are consistent with cooperation are also consistent with poor operational performance.

Although we attribute the performance differences that we measure to differences in the incentives of owned and independent regionals, one might question whether our results would not also be explained by differences in the skill levels of owned and independent regionals. Given that owned regionals tend to have higher labor costs, could they be employing more skilled employees – in particular, more skilled pilots? We believe that differences in skills are unlikely to be responsible for the performance effects that we find. It is hard to imagine how differences in the skill levels of pilots at owned and independent regionals would lead to differences in the performance of the majors using those regionals.³⁷ Specifically, even if owned regionals did have more skilled pilots, this might affect the frequency and/or severity of incidents and accidents at the regional but this is unlikely to affect departure delays at the major. One might hypothesize that pilot skill would affect arrival delays at the regional and, in turn, departure delays at the major. However, given that flight speed is determined by the airline (not the pilot) and the precise timing of a flight's arrival is determined by air traffic control (not the pilot), any possible relationship would seem to be quite weak.

VII. Conclusion

This paper has investigated whether – *at a given airport, on a given day* - the operational performance of majors using owned regionals differs from that of majors using independent regionals. Our results indicate that it does. We find that majors using only owned regionals at an airport experience delays that are about 3 minutes shorter, on average, than the delays experienced by majors using only independent regionals. This performance advantage increases on days with extreme rain and at congested times of day. We also find that integrated majors experience long delays and cancellations less frequently. These results are both statistically and economically significant. We find

³⁷ This is in contrast to willful actions of the pilots that can lead to longer flight delays, as evidenced by Lee and Rupp's (2007) finding that flight delays temporarily increased after airline pay cuts were announced.

much weaker evidence for an increase in the performance advantage of integrated airlines in snowy weather, but the two-staged least squares results on snow are at least broadly in line with those on rain and congestion. Overall, our findings are consistent with the view that integration improves a firm's ability to manage adaptation decisions.

We believe that this paper contributes to the existing literature in several ways. First, this paper is the first to show that, for similar transactions, the *operational* performance of integrated firms differs from the performance of non-integrated firms, and to measure the size of these performance differences. Second, we do this while attempting to control for the potential endogeneity of integration decisions both with numerous fixed effects and by instrumenting for the choice to vertically integrate. Third, our setting allows us to not only estimate performance differences but also shed light on their cause. In particular, the fact that airlines' ownership decisions are fixed in the short-run while the likelihood of adaptation decisions change on a daily basis provides us with a rich source of identification that is unavailable in many other settings.

References:

American Institute of Certified Public Accountants (2007), *Proposed AICPA Audit and Accounting Guide, Airlines*,
http://www.aicpa.org/download/exposure/Airline_Guide_ED.pdf

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

Ater, Itai (2009), “Internalizing Congestion at U.S. Hub Airports”, *mimeo*, Tel Aviv University.

Bajari, Patrick and Steven Tadelis (2001), “Incentives versus Transaction Costs: A Theory of Procurement Contracts,” *Rand Journal of Economics*, Autumn 2001, 32(3), 387-407.

Baker, George and Thomas Hubbard (2003), “Make versus Buy in Trucking: Asset Ownership, Job Design, and Information”, *American Economic Review* 93(3), 551-572.

Carlton, Dennis (1979), “Vertical Integration in Competitive Markets under Uncertainty”, *Journal of Industrial Economics* 27(3), 189-209.

Ciliberto, Federico (2006), “Does Organizational Form Affect Investment Decisions?” *Journal of Industrial Economics*, Volume LIV, No. 1, 63-93.

Federal Aviation Administration (1997), *Treatment of Values of Passenger Time in Air Travel*, <http://api.hq.faa.gov/economic/742SECT1.PDF>.

Forbes, Silke J. (2008), “The Effect of Air Traffic Delays on Airline Prices”, *International Journal of Industrial Organization* 26(5), 1218-1232.

Forbes, Silke J. and Mara Lederman (2009), “Adaptation and Vertical Integration in the Airline Industry”, *American Economic Review* 99(5), 1831-1849.

Forbes, Silke J. and Mara Lederman (2007), “The Role of Regional Airlines in the U.S. Airline Industry”, in Darin Lee (ed.), *Advances in Airline Economics, Volume 2: The Economics of Airline Institutions, Operations and Marketing*, Elsevier.

Forbes, Silke J. and Mara Lederman (2010), “Contract Form and Technology Adoption in a Network Industry”, *mimeo*, University of Toronto and UC San Diego.

Gibbons, Robert (2005), “Four Formal(izable) Theories of the Firm?” *Journal of Economic Behavior and Organization* 58, 200-245.

Gil, Ricard (forthcoming), “Revenue Sharing Distortions and Vertical Integration in the Movie Industry,” *Journal of Law, Economics and Organization*.

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.

Hirsch, Barry T. (2007), “Wage Determination in the U.S. Airline Industry: Union Power Under Product Market Constraints,” in Darin Leed (ed.), *Advances in Airline Economics, Volume 2: The Economics of Airline Institutions, Operations and Marketing*, Elsevier.

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

Hortacsu, Ali and Chad Syverson (2007), “Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices”, *Journal of Political Economy*, 115(2), 250-301.

Hubbard, Thomas (2001), “Contractual form and market thickness in trucking”, *Rand Journal of Economics* 32(2), 369-86.

Hubbard, Thomas (2008), “Viewpoint: Empirical Research on Firms’ Boundaries”, *Canadian Journal of Economics* 41(2).

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.

Lafontaine, Francine and Margaret Slade (2007), “Vertical Integration and Firm Boundaries: The Evidence”, *Journal of Economic Literature* 45(3).

Lee, Darin and Nicholas Rupp (2007), “Retracting a Gift: How Does Employee Effort Respond to Wage Reductions?”, *Journal of Labor Economics* 25(4), 725-762.

Levin, Jonathan and Steven Tadelis (forthcoming), “Contracting for Government Services: Theory and Evidence from U.S. Cities”, *Journal of Industrial Economics*.

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.

Mayer, Chris and Todd Sinai (2003), "Network Effects, Congestion Externalities, and Air Traffic Delays: Or Why All Delays Are Not Evil", *American Economic Review*, 93(4), 1194-1215.

Mazzeo, Michael J. (2003), "Competition and Service Quality in the U.S. Airline Industry", *Review of Industrial Organization*, 22(4), 275-296.

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

Morrison, S. and C. Winston (1989), "Enhancing the Performance of the Deregulated Air Transportation System", *Brookings Papers on Economic Activity*, 61-123.

Mullainathan, Sendhil and David Scharfstein (2001), "Do Firm Boundaries Matter?" *American Economic Review* 91(2), Papers and Proceedings of the Hundred Thirteenth Annual Meeting of the American Economic Association, 195-199.

Novak, Sharon and Scott Stern (2008), "How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry", *Management Science*, 54(12), 1963-1979.

Rupp, Nicholas G., Douglas H. Owens, and L. Wayne Plumly (2006), "Does Competition Influence Airline On-Time Performance?", in D. Lee (ed.), *Advances in Airline Economics*, Vol. I, Elsevier.

Rupp, Nicholas G. (2009), "Do Carriers Internalize Congestion Costs? Empirical Evidence on the Internalization Question", *Journal of Urban Economics*, 2009, 65, 24-37.

Tadelis, Steven (2002), "Complexity, Flexibility, and the Make-or-Buy Decision", *American Economic Review* 92, 433-37.

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.

Figure 1
Illustration of Identification Strategy

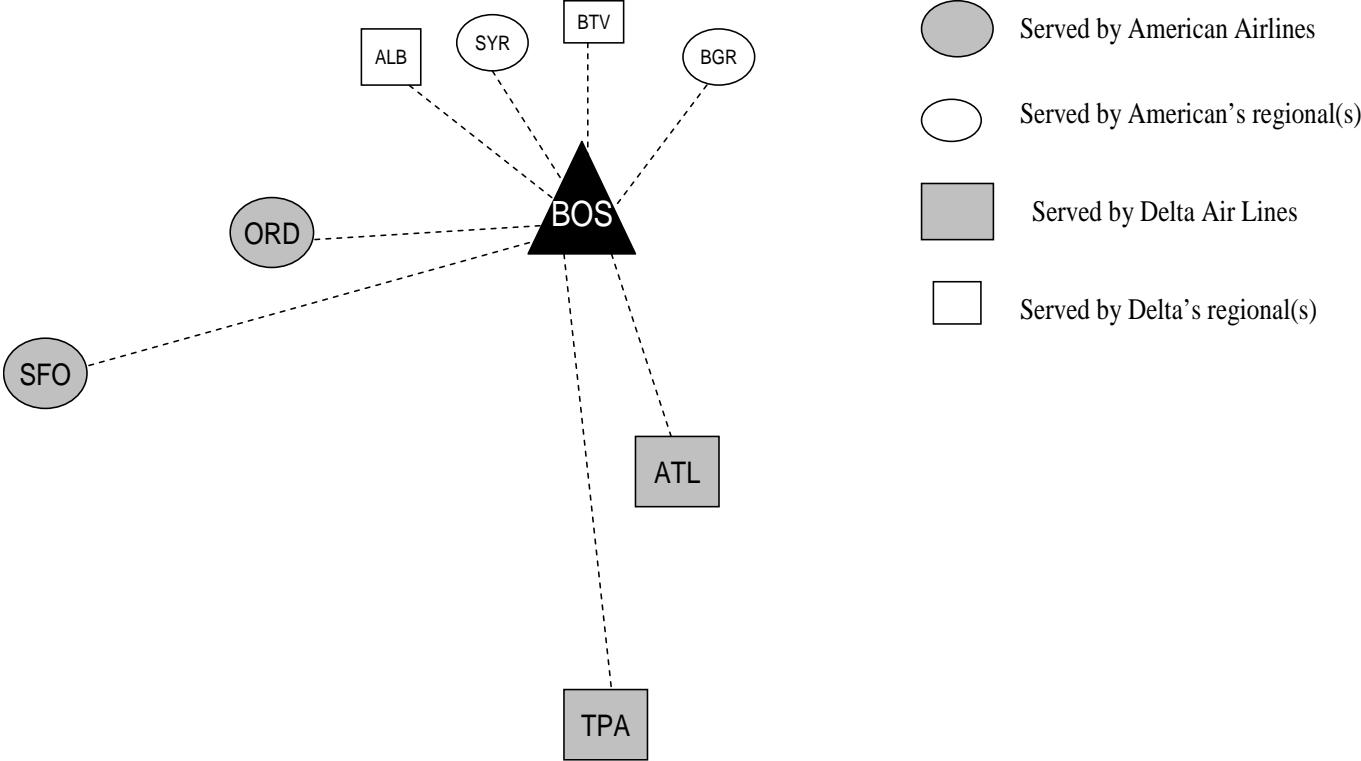


Table 1
Majors and Regional Partners in 2000
Regional carriers in bold are fully owned by the major

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 2
Summary of Delay Distribution, by Organizational Form

	Mean Delay	p25 Delay	p50 Delay	p75 Delay	Cancellations	N
<50% of Regional Flights at Airport are by Regional that is Owned	13.04	-3	0	12	0.05	790,749 (88% equal 0)
>50% of Regional Flights at Airport are by Regional that is Owned	9.88	-3	0	8	0.03	1,191,058 (71% equal 1)

Data is at the flight level. Table shows differences in delays and cancellations across flights operated by majors with more than/less than 50% of their regional flights at the departure airport of the flight operated by a regional that is owned.

Table 3
Variable Names and Definitions

Variable	Definition	Source
Dependent Variables		
<i>Departure Delay</i>	Difference between scheduled departure and actual departure of aircraft from the gate.	BTS On-time data
<i>Arrival Delay</i>	Difference between scheduled arrival and actual arrival of aircraft at the gate.	BTS On-time data
<i>Cancelled</i>	=1 if flight is cancelled	BTS On-time data
<i>Arrival Delay > 30 minutes or Cancelled</i>	=1 if arrival delay is more than 30 minutes or flight is cancelled	BTS On-time data
<i>Arrival Delay > 60 minutes or Cancelled</i>	=1 if arrival delay is more than 60 minutes or flight is cancelled	BTS On-time data
Ownership Variables		
<i>Fraction Owned Regional</i>	Fraction of major's regional flights at the departure airport that are operated by an owned regional partner	OAG & RAA data
<i>Owned Regional Flights</i>	Number of major's regional flights at the departure airport that are operated by an owned regional partner, in hundreds	OAG & RAA data
Weather and Congestion Variables		
<i>Rain</i>	Daily precipitation, on days with average temperature >32 degrees Fahrenheit (inches)	NOAA data
<i>Rain>95th Percentile</i>	=1 if rain at an airport on a day is greater than the 95 th percentile rain observed at that airport	NOAA data
<i>Snow</i>	Daily precipitation, on days with average temperature <=32 degrees Fahrenheit (inches)	NOAA data
<i>Snow>95th Percentile</i>	=1 if snow at an airport on a day is greater than the 95 th percentile snow observed at that airport	NOAA data
<i>Congestion</i>	Number of flights departing/arriving from/at an airport in an hour divided by maximum number of flights observed departing/arriving from/at that airport at any point in our sample	OAG data and Authors' construction
Airport Variables (defined for both departure and arrival airports)		
<i>Total Airport Flights</i>	Total number of domestic flights scheduled to depart from (arrive at) the airport on a day, in hundreds	OAG data
<i>Slot</i>	=1 if the airport is a slot-controlled airport (ORD, JFK, DCA); LGA is excluded, as described in text	Authors' construction
Airline-Airport Variables		
<i>Departs from Hub</i>	=1 if flight departs from an airport that is the operating airline's hub	Authors' construction
<i>Arrives at Hub</i>	=1 if flight arrives at an airport that is the operating airline's hub	Authors' construction
<i>Regional Flights</i>	A carrier's total number of <i>regional</i> flights at an airport on a day, in hundreds	OAG data

Table 4
Means of Selected Variables

	Mean	St Dev	Min	Max
Dependent Variables				
<i>Departure Delay (min)</i>	11.13	32.20	-15	1,076
<i>Arrival Delay (min)</i>	11.05	36.43	-1,298	1,076
<i>Cancelled</i>	0.04	0.19	0	1
<i>Arrival Delay > 30 minutes or Cancelled</i>	0.18	0.39	0	1
<i>Arrival Delay > 60 minutes or Cancelled</i>	0.11	0.31	0	1
Ownership Variables				
<i>Fraction Owned Regional</i>	0.56	0.45	0	1
<i>Owned Regional Flights (in hundreds)</i>	0.65	0.85	0	2.66
Weather and Congestion Variables (departure airports)				
<i>Rain (inches)</i>	0.11	0.35	0	12.56
<i>Rain Rain > 95th Percentile = 1</i>	1.29	0.79	0.02	12.56
<i>Rain > 0</i>	0.27	0.44	0	1
<i>Snow (inches)</i>	0.08	0.78	0	28.88
<i>Snow Snow > 95th Percentile = 1</i>	3.11	3.82	0.13	28.88
<i>Snow > 0</i>	0.03	0.18	0	1
<i>Congestion</i>	0.66			
Airline-Airport Variables				
<i>Departs from Hub</i>	0.63	0.48	0	1
<i>Arrives at Hub</i>	0.36	0.48	0	1
<i>Regional Flights (in hundreds)</i>	1.03	0.79	0.01	2.98

Table 5
Impact of Ownership on Delays
OLS Estimates

Dependent Variable	<i>Departure Delay (min)</i>					
	(5-1)	(5-2)	(5-3)	(5-4)	(5-5)	(5-6)
Ownership Variables						
<i>Fraction Owned</i>	-3.095** (0.213)	-3.108** (0.153)	-2.935** (0.151)	0.137 (0.172)		
<i>Fraction Owned*Rain>95th Percentile</i>			-4.239** (0.879)	-4.091** (0.816)	-4.128** (0.830)	
<i>Fraction Owned*Snow>95th Percentile</i>			3.207** (1.147)	3.204** (1.219)	2.543* (1.183)	
<i>Fraction Owned*Congestion</i>			-4.366** (0.458)	-1.944** (0.437)	-2.862** (0.437)	
<i>Owned Regional Flights</i>						-3.255** (0.208)
<i>Owned Regional Flights*Rain>95th Percentile</i>						-5.258** (1.172)
<i>Owned Regional Flights*Snow>95th Percentile</i>						4.835** (1.488)
<i>Owned Regional Flights*Congestion</i>						-1.941** (0.250)
Flight Level Controls						
<i>Congestion</i>	4.295** (0.416)	2.516** (0.204)	5.085** (0.340)	3.566** (0.328)	4.034** (0.331)	3.799** (0.244)
<i>Departs 8am-11am</i>	2.714** (0.118)	3.289** (0.102)	3.324** (0.102)	3.260** (0.102)	3.228** (0.103)	3.339** (0.102)
<i>Departs 11am-2pm</i>	5.727** (0.128)	6.316** (0.129)	6.302** (0.129)	6.358** (0.129)	6.368** (0.130)	6.277** (0.129)
<i>Departs 2pm-5pm</i>	9.880** (0.163)	10.760** (0.178)	10.798** (0.178)	10.818** (0.179)	10.802** (0.179)	10.798** (0.178)
<i>Departs 5pm-8pm</i>	13.539** (0.222)	14.583** (0.241)	14.578** (0.241)	14.569** (0.241)	14.577** (0.243)	14.555** (0.241)
<i>Departs 8pm-11pm</i>	12.733** (0.263)	13.643** (0.284)	13.687** (0.285)	13.663** (0.285)	13.627** (0.286)	13.607** (0.283)
<i>Departs 11pm-2am</i>	9.135** (0.558)	8.021** (0.497)	7.850** (0.489)	8.356** (0.487)	8.707** (0.486)	7.594** (0.470)
Airline-Airport Controls						
<i>Departs from Hub</i>	1.656** (0.387)	0.696** (0.265)	0.660* (0.261)	0.791** (0.263)		-0.550* (0.263)
<i>Regional Flights</i>	-0.857** (0.250)	0.755** (0.157)	0.660** (0.154)	0.030 (0.161)		3.369** (0.229)
<i>Arrives at Hub</i>	0.778** (0.132)	0.098 (0.093)	0.004 (0.093)	-0.480** (0.095)	-0.671** (0.093)	0.074 (0.092)
<i>Rain>95th Percentile*Departs from Hub</i>			-0.838 (1.759)	-0.830 (1.729)	-0.747 (1.706)	-2.537 (1.770)

<i>Snow>95th Percentile*Departs from Hub</i>	-1.594 (2.944)	-0.866 (2.979)	-1.055 (2.841)	0.960 (3.180)
<i>Rain>95th Percentile*Regional Flights</i>	2.230* (1.132)	2.298* (1.118)	2.227* (1.121)	6.488** (1.506)
<i>Snow>95th Percentile*Regional Flights</i>	2.140 (1.672)	2.396 (1.747)	2.399 (1.672)	-2.397 (2.333)
<i>Rain>95th Percentile*Arrives at Hub</i>	2.421** (0.450)	2.479** (0.447)	2.489** (0.443)	2.458** (0.450)
<i>Snow>95th Percentile*Arrives at Hub</i>	0.390 (0.550)	0.790 (0.552)	0.727 (0.549)	0.478 (0.549)

Departure Airport Controls

<i>Total Airport Flights</i>	0.445** (0.052)
<i>Slot Constrained</i>	1.802** (0.663)
<i>Rain>95th Percentile</i>	4.021** (1.509)
<i>Snow>95th Percentile</i>	4.938* (2.453)
<i>Rain>95th Percentile*Total Airport Flights</i>	0.856** (0.290)
<i>Snow>95th Percentile*Total Airport Flights</i>	0.754 (0.503)

Arrival Airport Controls

<i>Total Airport Flights</i>	0.293** (0.012)	0.275** (0.010)	0.281** (0.010)	0.271** (0.011)	0.284** (0.011)	0.266** (0.010)
<i>Slot Constrained</i>	0.896** (0.127)	0.787** (0.123)	0.842** (0.122)	-0.072 (0.124)	-0.350** (0.126)	1.044** (0.123)
<i>Rain>95th Percentile</i>	1.792** (0.355)	0.288 (0.243)	0.352 (0.244)	0.377 (0.243)	0.408+ (0.242)	0.351 (0.244)
<i>Snow>95th Percentile</i>	3.533** (0.778)	-0.541 (0.383)	-0.517 (0.381)	-0.439 (0.380)	-0.367 (0.377)	-0.502 (0.381)
<i>Rain>95th Percentile*Total Airport Flights</i>	0.731** (0.061)	0.585** (0.047)	0.407** (0.055)	0.401** (0.055)	0.393** (0.055)	0.407** (0.055)
<i>Snow>95th Percentile*Total Airport Flights</i>	0.585** (0.092)	0.477** (0.069)	0.450** (0.078)	0.420** (0.078)	0.404** (0.077)	0.450** (0.078)

Departure Airport-Date Fixed Effects		X	X	X	X	X
Airline Fixed Effects				X		X
Airline-Departure Airport Fixed Effects					X	
Observations	1,904,623	1,904,623	1,904,623	1,904,623	1,904,623	1,904,623
R Squared	0.04	0.16	0.16	0.17	0.17	0.16

Standard errors are clustered on airport-date. + significant at 10%; * significant at 5%; ** significant at 1%.

Table 6
Impact of Ownership on Delays
Instrumental Variables Estimates

Dependent Variable	<i>Departure Delay (min)</i>			
	(6-1)	(6-2)	(6-3)	(6-4)
Ownership Variables				
<i>Fraction Owned</i>	-2.812** (0.296)	-2.361** (0.306)	-2.391** (0.309)	
<i>Fraction Owned*Rain>95th Percentile</i>		-3.248* (1.555)	-3.202* (1.569)	-3.264* (1.454)
<i>Fraction Owned*Snow>95th Percentile</i>		-0.715 (2.009)	-0.381 (2.044)	-1.655 (2.018)
<i>Fraction Owned*Congestion</i>		-10.033** (0.960)	-10.114** (0.963)	-9.864** (0.986)
Flight Level Controls				
<i>Congestion</i>	2.526** (0.203)	8.433** (0.629)	8.483** (0.630)	8.129** (0.642)
Airline-Airport Controls				
<i>Departs from Hub</i>	0.724** (0.261)	0.592* (0.260)	0.611* (0.260)	
<i>Regional Flights</i>	0.626** (0.154)	0.666** (0.155)	0.666** (0.155)	
<i>Arrives at Hub</i>	-0.004 (0.097)	0.106 (0.098)	0.155 (0.098)	-0.610** (0.093)
<i>Mean (Rain>95th Percentile) across Major's Other Endpoints</i>			2.616** (0.887)	2.163* (0.873)
<i>Mean (Snow>95th Percentile) across Major's Other Endpoints</i>			3.645** (0.580)	3.393** (0.561)
<i>Mean (Rain>95th Percentile) across Regional's Endpoints</i>			-0.119 (0.837)	0.970 (0.843)
<i>Mean (Snow>95th Percentile) across Regional's Endpoints</i>			-0.577 (0.548)	-0.418 (0.524)
Departure Airport-Date Fixed Effects	X	X	X	X
Airline Fixed Effects				
Airline-Departure Airport Fixed Effects				X
Observations	1,904,623	1,904,623	1,902,866	1,902,866

Standard errors are clustered on airport-date. + significant at 10%; * significant at 5%; ** significant at 1%. All specifications estimated by two-stage least squares treating *Fraction Owned* and all of its interactions as endogenous. All specifications include departure time dummy variables. All specifications also include the additional airline-airport control variables that appear in Table 5 as well the arrival-airport control variables that appear in Table 5. The coefficients on these variables are not reported but are available upon request.

Table 7
Alternate Measure of Operational Performance

Dependent Variable	<i>=1 if Cancelled</i>	<i>Arrival Delay (min)</i>	<i>=1 if Cancelled or Arrival Delay>30min</i>	<i>=1 if Cancelled or Arrival Delay>60 min</i>
	(7-1)	(7-2)	(7-3)	(7-4)
Ownership Variables				
<i>Fraction Owned</i>	-0.017** (0.002)	-5.203** (0.355)	-0.042** (0.003)	-0.021** (0.003)
<i>Fraction Owned*Rain>95th Percentile</i>	-0.006 (0.008)	-2.676 (1.705)	-0.042** (0.015)	-0.030* (0.013)
<i>Fraction Owned *Snow>95th Percentile</i>	-0.019+ (0.012)	1.968 (2.229)	-0.009 (0.019)	-0.030+ (0.017)
<i>Fraction Owned*Congestion</i>	-0.019** (0.005)	-14.368** (1.124)	-0.134** (0.011)	-0.077** (0.008)
Airport-Date Fixed Effects	X	X	X	X
Observations	1,981,807	1,904,204	1,981,807	1,981,807

Standard errors are clustered on airport-date. + significant at 10%; * significant at 5%; ** significant at 1%. All specifications estimated by two-stage least squares treating *Fraction Owned* and all of its interactions as endogenous. All specifications include departure time dummy variables. All specifications also include the additional airline-airport control variables that appear in Table 5 as well the arrival-airport control variables that appear in Table 5. The coefficients on these variables are not reported but are available upon request. (7-2) has fewer observations because the delay variable is missing for flights that are cancelled.

Table 8
Alternate Weather Measures

Dependent Variable	<i>=1 if Cancelled or Arrival Delay>60min</i>	
	(8-1)	(8-2)
Ownership Variables		
<i>Fraction Owned</i>	-0.023** (0.003)	-0.020** (0.003)
<i>Fraction Owned*Rain</i>	-0.017* (0.008)	
<i>Fraction Owned*Snow</i>	-0.008+ (0.005)	
<i>Fraction Owned*Rain>0</i>		-0.007 (0.006)
<i>Fraction Owned*Snow>0</i>		-0.031* (0.014)
<i>Fraction Owned*Congestion</i>	-0.077** (0.008)	-0.078** (0.008)
Departure Airport-Date Fixed Effects	X	X
Observations	1,981,807	1,981,807

Standard errors are clustered on airport-date. + significant at 10%; * significant at 5%; ** significant at 1%. All specifications estimated by two-stage least squares treating *Fraction Owned* and all of its interactions as endogenous. All specifications also include the additional airline-airport control variables that appear in Table 5 as well the arrival-airport control variables that appear in Table 5. The coefficients on these variables are not reported but are available upon request. The coefficients on these variables are not reported.

Appendix A -Construction of Instruments

Fraction Owned measures the fraction of an airline's regional flights departing from a given airport that are operated by a regional that is owned. Our instruments for this variable are measures of the characteristics of the endpoint airports that are served from that airport by the regional carrier(s). For example, we instrument for Delta's value of *Fraction Owned* at the Boston airport with the characteristics of the endpoint airports that Delta's regionals serve from Boston. Our choice of instruments is motivated by the analysis in Forbes and Lederman (2009).

Specifically, for each flight that a major's regional operates from a particular airport, we calculate the following four characteristics of the arrival airport of that flight:

1. **Hub:** A dummy variable that equals one if the airport is a hub to the major.
2. **Precipitation:** The average annual precipitation at the airport. This average is taken over 25 years (1971-1995) of monthly weather data taken from the National Oceanic and Atmospheric Administration (NOAA). When precipitation is frozen (i.e. snow, hail, or freezing rain), this variable measures the water equivalent of the precipitation. Note that this is different from the depth of snowfall. The density of new snow is typically between 5% and 12% of water.
3. **Snowfall:** The average annual snowfall at the airport. This average is taken over 30 years (1971-2000) of annual snow data and reported by NOAA.
4. **# of Freezing Months:** The average number of months per year in which the average daily minimum temperature at the airport is below 32 degrees Fahrenheit. This average is taken over the 25 years of monthly weather data from NOAA.

After constructing these four measures for each regional flight from a particular airport, we calculate the average of these four variables over all of the flights that a particular major's regional(s) operate from a given airport during our sample period. For example, we would calculate the average of these four measures over all of the flights that Delta's regionals operate from Boston in the year 2000. This provides us four airline-airport level variables that we use as instruments for *Fraction Owned* (which is also an airline-airport level variable). We call these four variables *Fraction of Regional's Routes Arriving at Hub*, *Average Annual Precipitation at Endpoints Served by Regional*, *Average Annual Snowfall at Endpoints Served by Regional*, and *Average # of Months with Below Freezing Temperature at Endpoints Served by Regional*. The results of the first-stage regression of *Fraction Owned* on these variables, as well as the exogenous variables from the second-stage equation, are presented in Table 3.

In our sample, the correlation between average annual precipitation and daily rain is 0.022. The correlation between average annual precipitation and daily snow is 0.005. The correlation between average annual snowfall and daily snow is 0.004.

Appendix B – First Stage Regression

Dependent Variable	<i>Fraction Owned Regional</i>
Excluded Instruments	
<i>Fraction of Regional's Routes Arriving at Hub</i>	0.391** (0.005)
<i>Average Annual Precipitation at Endpoints Served by Regional</i>	0.007** (0.000)
<i>Average Annual Snowfall at Endpoints Served by Regional</i>	0.014** (0.000)
<i>Average # of Months with Below Freezing Temperature at Endpoints Served by Regional</i>	-0.286** (0.005)
Flight Level Control Variables	
<i>Congestion</i>	-0.049** (0.002)
Airline Controls	
<i>Departs from Hub</i>	0.214** (0.013)
<i>Regional Flights</i>	0.027** (0.009)
<i>Arrives at Hub</i>	-0.115** (0.002)
<hr/>	
Departure Airport-Date Fixed Effects	X
Observations	1,904,621
F-statistic on Excluded Instruments	3080.85
Prob>F	0.000
Total R-squared	0.5755
“Within” R-squared	0.1324

Standard errors are clustered on airline-airport-date. + significant at 10%; * significant at 5%; ** significant at 1%. All specifications include departure time dummy variables. All specifications also include the additional airline control variables that appear in Table 5 as well the arrival-airport control variables that appear in Table 4. The coefficients on these variables are not reported but are available upon request.