

HIAS-E-32

The Value of Relational Adaptation in Outsourcing: Evidence from the 2008 shock to the US Airline Industry

Ricard Gil

Johns Hopkins University

Hitotsubashi Institute for Advanced Study, Hitotsubashi University

Myongjim Kim

University of Oklahoma

Giorgio Zanarone

CUNEF

September 2016



Hitotsubashi Institute for Advanced Study, Hitotsubashi University
2-1, Naka, Kunitachi, Tokyo 186-8601, Japan
tel:+81 42 580 8604 <http://hias.ad.hit-u.ac.jp/>

HIAS discussion papers can be downloaded without charge from:
<http://hdl.handle.net/10086/27202>
<https://ideas.repec.org/s/hit/hiasdp.html>

All rights reserved.

The Value of Relational Adaptation in Outsourcing: Evidence from the 2008 shock to the US Airline Industry^{*}

Ricard Gil	Myongjin Kim	Giorgio Zanarone
Johns Hopkins U Hitotsubashi Institute for Advanced Study Hitotsubashi University	University of Oklahoma	CUNEF

September 2016

Abstract

In the airline industry, ex-post adaptation of flight schedules is necessary in the presence of bad weather conditions. When major carriers contract with independent regionals, conflicts over these adaptation decisions typically arise. Moreover, the celerity of needed adjustments requires that adaptation be informal, and hence enforced relationally. In this paper, we theoretically analyze, and empirically test for, the importance of relational adaptation in the airline industry. Our model shows that for relational contracts to be self-enforcing, the long-term value of the relationship between a major and a regional airline must be at least as large as the regional's cost of adapting flight schedules across joint routes. Thus, when facing a shock that forces it to terminate some routes, the major is more likely to preserve routes outsourced to regional airlines that have higher adaptation costs, as the value of the major's relationship with those regionals is larger. We analyze the evolution of U.S. airline networks around the 2008 financial crisis, and we find that consistent with our theoretical predictions, regional routes belonging to networks with worse average weather, and hence higher adaptation costs, were more likely to survive after the shock.

Keywords: Relational contracting, adaptation, natural experiment, airlines, outsourcing

JEL codes: L14, L22, L24, L93

^{*} We thank Ricardo Alonso, Benito Arruñada, Mikhail Drugov, John de Figueiredo, William Fuchs, Bob Gibbons, Bentley MacLeod, Joanne Oxley, Mike Powell, Nick Vikander, and audiences at Aarhus University, Pompeu Fabra University, Tsinghua University, Johns Hopkins Carey, Hitotsubashi University, the "Natural Experiments and Controlled Field Studies" Workshop in Ohlstadt (Germany), and the ALEA and SIOE conferences, for useful comments. This study received financial support from the Spanish Ministry of Science and Education through grant ECO2014-57131-R. The usual disclaimer applies.

1. Introduction

Adaptation to change is paramount to the success of organizations and markets. As Hayek (1945: 523) states, “economic problems arise always and only in consequence of change”. Reinforcing Hayek’s point, Williamson (1991: 278) argues that “adaptability is the central problem of economic organization,” and warns that coordinated responses to change, which may require the cooperation of multiple parties with diverging interests and an urgent implementation, can rarely be achieved in spot markets.

While agreeing on the importance of this fundamental adaptation problem and on the difficulty to address it through spot market exchange, economic theory has emphasized at least three different approaches to finding a solution. First, allocating decision authority to a “boss”, as in firms (e.g., Simon 1951; Williamson 1991; Hart and Holmstrom 2010). Second, agreeing *ex ante* on formal procedures that facilitate the renegotiation of key decisions, as in procurement and construction contracts (Bajari and Tadelis 2001; Chakravarty and MacLeod 2009). Third, using relational contracts—informal agreements where coordinated adaptation is ensured by the parties’ desire to maintain a long-term collaborative relationship (Baker, Gibbons and Murphy 2011).

This paper sheds light onto the incidence and relevance of the third solution, relational adaptation, by empirically assessing the importance of relational contracts as a solution to the adaptation problem in outsourcing relationships. We exploit a setting that is particularly well suited to study relational adaptation—namely, the networks of major and regional airlines in the US. To begin with, as shown in previous empirical work by Forbes and Lederman (2009), adaptation is key in this industry: when major carriers rely

on outsourcing agreements with independent regional carriers to serve local routes, important coordination challenges arise, as adverse weather conditions and other unexpected contingencies require adapting flight schedules in ways that may collide with the regionals' profit-maximization objectives. In addition, outsourcing regional transportation offers significant labor cost advantages to major airlines, as regional airlines are scarcely unionized. Thus, while allocating decision authority to major airlines through integration is a potential solution to the adaptation problem (see Forbes and Lederman, 2009, 2010, for convincing evidence on the incidence of adaptation in airlines), it is not always a feasible or desirable one. At the same time, formal contracts between major and regional airlines are inherently incomplete, both *ex ante* and *ex post*, when it comes to adaptation. On one hand, it is difficult to foresee *ex ante* under what circumstances flights on a specific route should be rescheduled, and on the other hand, rescheduling decisions must be implemented quickly, and full cooperation by the regional airlines may be hard to verify. Finally, U.S. major airlines and independent regional airlines are typically engaged in long-term business relationships, and hence are likely to secure the more informal aspects of their collaboration, such as the adaptation of flight schedules, through relational contracts.

1.1. Overview of the results

To study relational adaptation in the U.S. airline industry, we begin by developing a simple model, adapted from Baker *et al.* (2011), where a major airline promises quasi-rents to an independent regional partner in exchange for compliance with the major's schedule adaptation decisions. Following Levin (2002), we assume that deviations from this relational contract on one route trigger termination of the relationship between the

major and the regional on all routes, as that is the worst credible punishment. A key point is that given the lack of court-enforcement, relational adaptation is achieved if, and only if the long-term value of the relationship between the major and the regional is at least as large as the regional's present adaptation costs across all the routes in their joint network. Thus, if the relational enforcement constraint is binding, the relationship's value must be larger the larger the network-wide adaptation cost. A testable implication of this fact is that facing a negative shock that induces the major to cut some routes, and given that cutting routes may jeopardize the network's viability, the major will choose to preserve those routes that belong to networks with high aggregate adaptation costs, and hence high long-term value.

We test our prediction on a comprehensive dataset of relationships between U.S. major and regional airlines. As an exogenous industry-wide shock that induces major airlines to cut routes, we exploit the economic crisis that unraveled after the collapse of Lehman Brothers in 2008. Following Forbes and Lederman (2009), we proxy adaptation costs within a major-regional network by the extent of bad weather conditions (precipitation, snow, and the lack of clear skies as inversely measured by the number of freezing months) across the routes in the network—the rationale being that in the presence of adverse weather, major airlines are forced to reschedule more regional flights, which in turn increases the regional airlines' adaptation costs. Because a given regional typically operates different sets of routes for different majors and a major may use different regionals to operate the same route, we can include major airline, regional airline, and route fixed effects in our regressions, thus controlling for the possibly endogenous assignment of routes to major and regional carriers.

Our empirical results support the hypothesis that the long-term relationships between major and regional airlines are used to informally enforce the efficient adaptation of flight schedules. Controlling for a rich set of route and network characteristics, as well as for major airline and regional airline fixed effects, we find that outsourcing routes in networks characterized by high snow and precipitation are more likely to survive, and less likely to see their number of daily flights reduced, following the Lehman Brothers shock in 2008. In contrast, routes in networks characterized by frequent freezing weather—which Forbes and Lederman (2009) suggest to be a proxy for clear skies—are less likely to survive and more likely to experience flight reductions. Our results are robust to the inclusion of route-level weather, route fixed effects, and route-major and route-regional fixed effects.

As a placebo test, we also study route survival in the 2003-2006 period, and we find that absent a negative shock that pushes the major airlines to reduce route portfolios in the less valuable networks, routes in bad weather major-regional networks are as likely to be cut or downsized as routes in good weather major-regional networks. We also study route survival around a different industry-wide shock—namely, the shock following the 9/11 terrorist attacks in 2001. The results there also imply overall no connection between network-level weather variables and the likelihood of survival of outsourcing routes. This is not surprising for two reasons. First, the impact of the 9/11 shock on the U.S. airline industry, and on the U.S. economy as a whole, was shorter and more transitory than the impact of the Lehman Brothers shock. In the same way, the shock itself did not structurally affect the US airline industry in the same way as the 2008 financial crisis did. Second, the use of regional airlines was sparse in 2001 relative to 2008. Therefore it is

not surprising to find that US major airlines did not significantly adjust their outsourced regional networks in the 1999-2003 period.

We conclude our empirical investigation by further exploring the strategic decisions of major airlines when they decide to sever an outsourcing relationship on a route. The major airline may decide to stop flying that route completely or adjust to the shock through alternative margins—namely, vertical integration and route reallocation to other regionals. We find that routes were more likely to be integrated after the Lehman Brothers shock if formerly flown by a regional with better network-wide weather, and hence lower continuation value. Similarly, when majors reallocated routes to other partners, they exclusively relied on existing partners—that is, regional airlines that were operating other routes for the same major airline before the shock. That means that following the 2008 shock, major airlines cut some partnerships while intensifying their relationships with the stronger existing partners.

1.2. Contribution to the literature

1.2.1. Adaptation

Early empirical works have focused on how formal price adjustment provisions facilitate adaptation and reduce its costs in procurement contracts. In particular, Masten and Crocker (1985) show that gas supply contracts stipulate lower penalties against breach when the supplier can more easily store the unsold gas or sell it to alternative clients—that is, when breach by the buyer in the face of unforeseen market conditions is more likely to be efficient. Crocker and Reynolds (1993) show that aircraft engine procurement contracts contain more complete pricing provisions when intertemporal and

technological uncertainties are low (reduced need for ex post adaptation) and when the supplier has a history of legal disputes with clients (high cost of negotiated adaptation).

More recent studies have focused on the allocation of authority as a means to facilitate adaptation. Arruñada et al. (2001), and Zanarone (2009, 2013), show that automobile distribution agreements assign to car manufacturers the right to adapt the dealers' performance and service standards ex post in networks where dealers are more likely to freeride on the brand. Forbes and Lederman (2009, 2010) show that U.S. major airlines vertically integrate into regional transportation in bad weather routes, where there is more need for adapting flight schedules, and that vertical integration reduces the flight delays and cancellations that would arise in the absence of coordinated adaptation.

Unlike our paper, none of the empirical studies discussed above focuses on relational contracts—in the sense of informal, self-enforcing agreements—as mechanisms to facilitate adaptation. The only other empirical paper we are aware of that explores the link between relational contracts and adaptation is Barron et al. (2015). They study contracts between a movie exhibitor and multiple distributors, and show that for any given movie, the exhibitor's revenue share is adjusted upwards if the exhibitor continues showing the movie. They also show that price adjustments are more generous when the exhibitor's opportunity cost of showing the movie is high. Since price adjustments are not prescribed by the formal contracts signed by the parties at the beginning of their relationship, these findings suggest that the adjustments may serve as implicit bonuses to reward the exhibitor for adapting movie schedules.

Our empirical evidence on relational adaptation differs from, and complements that in Barron et al. (2015). While they look in great detail at how a given relational contract is

adapted over time holding the set and depth of relationships constant, we provide evidence on the economic significance of relational adaptation contracts by showing that those with higher long-term value are more likely to survive market shocks.

1.2.2. Relational contracts

Economic theory has extensively investigated relational contracts—that is, contracts that are too rooted in the parties’ relationship to be verifiable by courts, and hence must be self-enforcing (see MacLeod 2007, and Malcomson 2013, for up-to-date reviews of the theoretical literature). The predictions of relational contracting theories have been confirmed empirically by both case studies (e.g., Macauley 1963; Fast and Berg 1975; Foss 2003; Helper and Henderson 2014) and econometric evidence (Gil and Zananone, 2015 and 2016, offer an up-to-date review and critical assessment of the empirical literature). As mentioned above, none of the existing empirical works, with the exception of Barron et al. (2015), study relational contracts as a means to achieve adaptation to unforeseen contingencies.

Methodologically, the empirical paper on relational contracting that is perhaps most closely related to ours is Macchiavello and Morjaria (2015). In their study of flower export agreements, they propose the idea, which we exploit in this paper, that the long-term value of a relational contract may be estimated by measuring the largest reneging temptation, conditional on the parties being in a relationship. Macchiavello and Morjaria (2015) use this idea to show that informal contracts between Kenyan flower exporters and their clients reduce the volume of stipulated flower deliveries following unexpected increases in the spot market price, so that the exporters’ reneging temptation, given by the product of spot price and relational quantity, remains constant, and hence equal to the

relationship's long-term value. In contrast, our empirical exercise shows that positive cross-sectional variations in the long-term value of relationships, as measured by the present reneging temptation, increase the parties' willingness to preserve those relationships in the face of a shock. A second, important difference between our study and Macchiavello and Morjaria (2015) is the type of relational contract under study. While they focus on simple quantity-price agreements in an institutional environment characterized by weak court enforcement, we focus on complex outsourcing agreements in a strong institutional environment, and show that despite the presence of efficient courts, because of the complexity of these agreements, major and regional airlines rely on relational contracts to adapt them to unforeseen events.

The paper is organized as follows. While section 2 describes the US airline industry, section 3 presents an illustrative model with testable implications that we take to data. In section 4 we describe our data. Section 5 presents our empirical methodology and the main results of our empirical analysis. Section 6 presents our placebo test and our analysis of the 9/11 shock, and examines alternative margins of network adjustment—vertical integration and route reallocation—employed by major airlines in response to the 2008 shock. Section 7 concludes.

2. Adaptation in the U.S. airline industry

Major airlines fly routes using either their own fleets or those of regional airlines. Regional airlines may be owned by the major airlines or independently owned. Major airlines invest considerable planning efforts into designing schedules of flights and landing times for their regional networks, so that passengers can reach their destinations

in a timely manner. While in regular and good weather adjustments to the landing schedule are not necessary, they do become necessary under bad weather conditions. Because landing in bad weather takes longer and requires more caution, the number of landing slots available is reduced. Specifically, airport authorities unilaterally decrease the number of slots through Ground Delay Programs (GDPs hereafter), and they do so proportionally given each airline's original schedule. Major airlines then rearrange their schedule using the slots available to them and their integrated regionals, but they do not have control (in the sense of real authority) over slots assigned to independently owned regionals that they outsource to.

During GDPs, airlines are able to exchange slots with their independent regionals through a formal mechanism called SCS (Slot Credit Substitution). SCS is a centralized system where the major airline asks for an immediate time slot from any airline and often from its independent regional, in exchange for a later time slot (Schummer and Vohray, 2013; Vossen and Ball, 2006). If the regional accepts, it foregoes a landing slot and thus it has to delay or even cancel one of its flights. Three important features of this process, which are captured by our theoretical model in the next section, should be stressed. First, there is a conflict of interest between majors and regionals concerning the adaptation of flight schedules. While majors are residual claimants of flight revenues and pay a fixed fee per flight to the regionals, delaying flights requires the regional to distort their employees' schedules, which may result in higher labor and logistics costs (Forbes and Lederman 2009).

Second, since independent regionals own, and hence control their slots, these conflicts cannot be easily resolved via formal contracts. On the one hand, maximization

of network-wide profits requires that flights be rescheduled quickly in the presence of bad weather, so “haggling” between the major and the regional may be costly (e.g., Hart and Moore 2008). On the other hand, allocating decision rights *ex ante* may also fail to elicit efficient adaptation. While contracts often assign to majors the formal right to reschedule their regionals’ flights, as owners of their slots the regionals have an option to refuse, perhaps in the hope to bargain harder with the major, or just to save on labor costs. The major would then have to sue the regional for contract breach, and given the environmental changes due to bad weather conditions, a court may well “complete” the contract in favor of the regional by applying flexible legal doctrines (e.g., Schwartz 1992). Anticipating that, the major may have an incentive to breach. Based on these considerations, and following the theoretical approach in Baker et al. (2011) and Barron et al. (2015), we assume in our model that rescheduling decisions are non-contractible.

Finally, it appears that despite the conflict of interests and contractual frictions discussed above, independent regionals do routinely accept the majors’ schedule adjustment decisions, and major airlines compensate them for doing so. In particular, while we do not have evidence of *ex-post* monetary payments being made by the majors in exchange for the adjustments, conversations with industry practitioners suggest that majors informally count the flights cancelled by their regional partners as a consequence of requested slot exchanges as valid for the yearly minimum number of flights that based on the outsourcing agreements, the regionals have to reach in order to receive the fixed operation fee. In other words, the major’s choice not to enforce the minimum clause in the contract *ex-post* serves as an informal performance “bonus”.

Considering this institutional background and the demand for relational adaptation between major and regional airlines, we formally model in the next section relational contracting as a solution to the adaptation problem in this industry.

3. A simple model of relational adaptation in the airline industry

In this section we present a simple model of relational adaptation that captures the key features of the U.S. airline industry as described above, and allows us to generate empirically testable predictions.

3.1. Model setup

There are a major airline, M, and an independent regional airline, R, which may operate up to N routes on M's behalf. Both M and R are risk-neutral, live forever, and discount next-period payoffs by the factor $\delta \in [0,1]$. Time evolves in discrete periods.

At stage zero of any given period t , M and R observe a state, $z_t \in \{0,1\}$, whose realization is independent across periods. In "normal times", $z_t = 0$, and R's cost of operating route i in the current period is given by \underline{k}_i . In the presence of an unlikely negative shock, $z_t = 1$, and R's operating cost rises to $\underline{k}_i + \varepsilon_{it}$, where ε_{it} is a random adjustment, also independent across periods, distributed along the \mathbb{R}^+ line with pdf $f(\cdot)$, and cdf $F(\cdot)$. To capture the idea that the shock is unlikely, we assume $Pr(z_t = 0) \approx 1$, and $Pr(z_t = 1) \approx 0$, for all t .

At stage one, after observing the shock state, M decides which of the N routes to outsource to R. We write $h_{it} = 1$ if route i is outsourced to R, and $h_{it} = 0$ otherwise. If $h_{it} = 1$, M offers to pay to R a fixed fee, $r_{it} \in \mathbb{R}$, in exchange for operating the route. If R accepts M's offer, R incurs the operating cost, and the game moves to stage two. If $h_{it} = 0$, or if R rejects M's offer, M receives payoff m_i^0 , R receives payoff zero, and the game moves to period $t + 1$. We may interpret m_i^0 as the maximum between M's payoff from not serving the route and M's payoff from operating the route with its own planes or by using a vertically integrated regional company. If $h_{it} = 0$ and period t has been normal ($z_t = 0$), in period $t + 1$ the game repeats itself identically. In contrast, if $h_{it} = 0$ and there has been a negative shock in period t ($z_t = 1$), the game ends from period $t + 1$ and thereafter with probability $q_i \in [0,1]$, in which case M receives m_j^0 , and R zero, on every route $j \in N$, that is, including on those routes that had been outsourced to R at time t .

By assuming that lack of outsourcing on one route may trigger termination of the whole regional network, we capture the idea that in the presence of a shock, restructuring of M's portfolio of outsourced routes may cause negative network-wide externalities. For instance, R may experience temporary financial constraints due to the shock, such that if route i is not outsourced, and hence R does not receive the operation fee r_i , R may go bankrupt and be forced to exit the market.

At stage two of period t , after the shock state z_t has been realized and the outsourcing decisions have been made, M and R observe the route state $w_{it} \in \{0,1\}$, where $w_{it} = 1$ denotes bad weather affecting flights on route i , $w_{it} = 0$ denotes good weather, the

probability of bad weather is $p_i \in [0,1]$, the probability of good weather is $1 - p_i$, and we assume for simplicity that weather states are independent across routes and time periods.

At stage three, after observing w_{it} , R chooses the adaptation decision, $d_{it} \in \{0,1\}$, at cost $d_{it}c_i$. Consistent with the features of the US airline industry described in section 2, we say that adaptation occurs, that is, $d_{it} = 1$, if R gives up one of its slots on route i to M, thereby rescheduling and potentially delaying its own flights on that route. Conversely, $d_{it} = 0$ if R does not give slots to M, and hence R does not need to reschedule its flights on route i . The adaptation cost, $c_i > 0$, may include the workers' extra hours and additional maintenance costs that R must incur if its flights on route i are delayed as a consequence of giving a slot to M. If $d_{it} = 1$, M may pay a bonus, $b_{it} \in \mathbb{R}$, to compensate R's adaptation cost.

Finally, at stage four, M receives profit $m_i(d_{it}, w_{it})$ from each outsourced route i given the realized weather and R's adaptation decision.

We make the following two assumptions:

A1: $m_i(\mathbf{1}, \mathbf{1}) - c_i > m_i(\mathbf{0}, \mathbf{1})$ and $m_i(\mathbf{0}, \mathbf{0}) > m_i(\mathbf{1}, \mathbf{0}) - c_i$, for all i .

A2: $p_i m_i(\mathbf{0}, \mathbf{1}) + (1 - p_i) m_i(\mathbf{0}, \mathbf{0}) < m_i^0$, for all i .

Assumption A1 implies that it is efficient to reschedule flights only in the presence of bad weather. This creates a potential conflict of interest between M and R, as adaptation benefits M but is costly for R. Assumption A2 implies that in the absence of efficient adaptation, it is optimal for M not to outsource a route. We interpret this assumption as the joint result of the intense competition M may face from other airlines, and to the well

documented fact that major airlines vertically integrate into poorly performing routes (Forbes and Lederman 2009).

It is useful to write M's and R's gross expected payoffs and the gross expected surplus in period t , conditional on the outsourcing decision (h_{1t}, \dots, h_{nt}) and on efficient adaptation, respectively, as:

$$\pi_M(h_{1t}, \dots, h_{nt}) \equiv \sum_i h_{it} [p_i m_i(1,1) + (1 - p_i) m_i(0,0)] + \sum_i (1 - h_{it}) m_i^0, \quad (1)$$

$$\pi_R(h_{1t}, \dots, h_{nt}) \equiv - \sum_i h_{it} [k_i + p_i c_i], \text{ and} \quad (2)$$

$$S(h_{1t}, \dots, h_{nt}) \equiv \pi_M(h_{1t}, \dots, h_{nt}) + \pi_R(h_{1t}, \dots, h_{nt}). \quad (3)$$

We make the following informational assumptions.

A3: w_{it} , $m_i(d_{it}, w_{it})$, c_i are observable but non-verifiable, for all i and t .

A4: d_{it} is observable but non-verifiable, for all i and t .

Assumption A3 is standard in the incomplete contracting literature, and it implies that efficient flight adaptation is ex ante non-contractible (e.g., Grossman and Hart 1986; Hart and Moore 1988). Assumption A4 implies that efficient flight adaptation decisions are formally non-contractible even ex post, after weather is observed (e.g., Baker et al. 2011). This is consistent with the institutional features of the airline industry discussed above, according to which rescheduling decisions are too urgent, fast, and state-contingent to be formally contracted at a reasonable cost, either ex ante or ex post (Forbes and Lederman 2009).

3.2. Spot market contracts

Suppose M and R rely on formal, spot market contracts to govern their outsourcing agreement. Since adaptation decisions and contingent bonuses are non-contractible, at stage 3 of any given period t , M will pay no bonus to R: $b_{it} = 0$ for all i . Anticipating that, R will not adapt flight schedules, irrespective of the realized weather state: $d_{it} = 0$ for all i . But then, our assumption A2 implies that at stage one, R will not outsource any routes to R, irrespective of the realized shock state: $h_{it} = 0$ for all i . As a result, M's expected per period payoff will be $\pi_M(0, \dots, 0)$, R's payoff will be zero, and the total surplus will be $S(0, \dots, 0) = \pi_M(0, \dots, 0)$.

3.3. Relational adaptation contracts

While rescheduling decisions are formally non-contractible, M and R may still improve on the spot market by entering a relational adaptation contract, whereby R promises to execute the efficient state-contingent decision schedule, $d_{it}^*(w_{it}) \equiv w_{it}$, on all the outsourced routes and in all periods, in exchange for the quasi-rents from continuation of the relationship with M.

Formally, a relational adaptation contract is a complete plan for the relationship between M and R, which specifies, for any realized history of play up to any given period: M's outsourcing decisions and upfront operation fees as a function of the state (normal times vs. shock), R's adaptation decisions as a function of weather, and the discretionary bonuses M has to pay R conditional on R's adaptation decisions. We say that a relational adaptation contract is self-enforcing if it describes a subgame perfect equilibrium of the repeated game between M and R. Following Levin (2002), we assume

that if M and R enter a relational adaptation contract, deviations on one route (that is, an unexpected outsourcing decision or upfront payment, R's failure to reschedule flights in the presence of bad weather, or M's failure to pay the bonus after R reschedules) are punished through reversion to spot market contracting on all the outsourced routes, as that is the worst credible punishment. Given perfect public monitoring and the absence of liquidity constraints, the optimal relational contract is stationary, in the sense that conditional on the realized states, z_t and w_{it} , outsourcing and adaptation decisions and payments on any given route i are the same in every period (MacLeod and Malcomson 1989; Levin 2003). Accordingly, we hereafter drop the time subscripts from all equations.

Notice that since the negative shock is unlikely, it does not affect M's and R's expected payoffs from the relationship. Thus, we can separately analyze relational adaptation in normal times (that is, conditional on $z = 0$) and in the presence of a shock ($z = 1$).

3.3.1. Normal times

Given assumption A2, M's optimal relational contract in normal times can be characterized as a vector of stationary outsourcing decisions, (h_1^*, \dots, h_n^*) , such that $h_i^* = 1$ if $d_i \equiv w_i$, and $h_i^* = 0$ otherwise, and a set of payments that guarantee M's and R's participation and efficient adaptation on the outsourced routes. Formally, M solves:

$$\max_{h_i, r_i, b_i} \{ \pi_M(h_1, \dots, h_n) - \sum_i h_i (p_i b_i + r_i) \},$$

subject to the following constraints:

$$\pi_M(h_1, \dots, h_n) - \sum_i h_i (p_i b_i + r_i) \geq \pi_M(0, \dots, 0), \quad (4)$$

$$\pi_R(h_1, \dots, h_n) + \sum_i h_i (r_i + p_i b_i) \geq 0, \quad (5)$$

$$\sum_i h_i (b_i - c_i) + \frac{\delta}{1-\delta} [\pi_R(h_1, \dots, h_n) + \sum_i h_i (r_i + p_i b_i)] \geq 0, \text{ and} \quad (6)$$

$$-\sum_i h_i b_i + \frac{\delta}{1-\delta} [\pi_M(h_1, \dots, h_n) - \sum_i h_i (p_i b_i + r_i)] \geq \frac{\delta}{1-\delta} \pi_M(0, \dots, 0). \quad (7)$$

Conditions (4) and (5) are M's and R's participation constraints, respectively. Conditions (6) and (7) are R's and M's incentive constraints, which ensure, respectively, that R be willing to reschedule flights following bad weather (condition 6), and M be willing to pay the promised contingent bonuses (condition 7), in the highest-temptation state—that is, in case of bad weather on all the outsourced routes.¹

Summing up (6) and (7), we obtain a necessary condition for the relational adaptation contract to be self-enforcing:

$$\sum_i h_i c_i \leq \frac{\delta}{1-\delta} [S(h_1, \dots, h_n) - S(0, \dots, 0)]. \quad (8)$$

In fact, it is easy to check that condition (8) is also sufficient, in the sense that there are operation fees and bonuses such that if (8) holds, (4) through (7) hold as well, and M extracts the whole surplus, $S(h_1, \dots, h_n)$.

If condition (8) is not satisfied—either because R's cost of adapting schedules, on the left-hand side, is too large, or because the present expected value of the relationship, on the right-hand side, is too small—efficient adaptation on the agreed routes cannot occur, so M will need to outsource fewer routes to R in order to keep the relational adaptation contract within its self-enforcing range.

Let R's total adaptation cost in the worst weather state, and the net continuation value of the relationship between M and R, under the optimal outsourcing decision, be $\sum_i h_i^* c_i$

¹ We omit the constraint that R be willing to accept the contingent bonus if negative because this constraint is looser than (6), and hence redundant.

and $V^* \equiv \frac{\delta}{1-\delta} [S(h_1^*, \dots, h_n^*) - S(0, \dots, 0)]$, respectively. Then, condition (8) implies that the total adaptation cost, $\sum_i h_i^* c_i$, is a lower bound for the value of the relationship, V^* (Macchiavello and Morjaria 2015). In particular, when the discount rate is small enough, so that (8) is binding, the relationship's value must equal the adaptation cost: $V^* = \sum_i h_i^* c_i$.

3.3.2. A negative shock

In the unlikely event that a negative shock occurs (that is, if $z = 1$), its impact on the outsourcing decision is twofold. On the one hand, R's present cost of operating a given route i is raised to $\underline{k}_i + z\varepsilon_i$, which disfavors outsourcing. On the other hand, if the route is not outsourced to R, the network terminates with probability q_i in the following period. Thus, routes that are optimally outsourced in normal times may no longer be outsourced following a shock, depending on the balance between the present costs and future benefits of outsourcing those routes, as defined above.

Assume route i is outsourced in normal times ($h_i^* = 1$), and let M's per period expected payoff from outsourcing route i in normal times be:

$$S_i \equiv p_i[m_i(1,1) - c_i] + (1 - p_i)m_i(0,0) - \underline{k}_i.^2$$

Given the above definitions and assumptions, M will continue to outsource route i to R following a shock if, and only if:

$$S(h_1^*, \dots, h_n^*) - \sum_{j=1}^n h_j^* \varepsilon_j + \frac{\delta}{1-\delta} S(h_1^*, \dots, h_n^*) > S(h_1^*, \dots, h_n^*) - S_i - [(\sum_{j=1}^n h_j^* \varepsilon_j) - \varepsilon_i] + q_i \frac{\delta}{1-\delta} S(0, \dots, 0) + (1 - q_i) \frac{\delta}{1-\delta} S(h_1^*, \dots, h_n^*), \text{ or}$$

² Recall that under an optimal relational adaptation contract, M extracts all the expected surplus from the outsourced routes.

$$q_i V^* + S_i > \varepsilon_i.$$

If self-enforcement condition (8) is binding, so that $V^* = \sum_i h_i^* c_i$, then the above optimal outsourcing decision rule can be rewritten as:

$$q_i \sum_i h_i^* c_i + S_i > \varepsilon_i. \quad (9)$$

Thus, we can define the probability that a given route i , which is optimally outsourced in normal times, is also outsourced following a shock—that is, the probability that route i survives the shock—as:

$$Pr(Survival_i) \equiv F(q_i \sum_i h_i^* c_i + S_i). \quad (10)$$

Condition (10) implies that the likelihood of survival of an outsourced route i after the shock, $Pr(Survival_i)$, is higher the higher R's aggregate adaptation cost across all the routes that were optimally outsourced prior to the shock—that is, $Pr(Survival_i)$ increases in $\sum_i h_i^* c_i$.³

Proposition: *Assume the self-enforcement constraint in the relational adaptation contract, (8), is binding. Then, the likelihood that R continues to outsource a given route i to R following a shock increases in $\sum_i h_i^* c_i$, R's aggregate adaptation cost across all the routes that M optimally outsourced to R before the shock.*

The intuition behind our Proposition is simple. If M and R use a relational contract to efficiently adapt flight schedules to changes in weather, the binding relational

³ Our assumption A2 matters here, because it rules out the possibility that M may outsource a route to R without requiring efficient adaptation—that is, it rules out the possibility that $h_i^* = 1$ and $d_i(1) = 0$ for some route i . If M tolerated inefficient adaptation on some routes, $\sum_i h_i^* c_i$ may overestimate the relationship's long-term value. Again, we find our assumption consistent with the intense competition that characterizes the U.S. airline industry, and on the existence of vertical integration as a way, albeit a costly one, to overcome poor adaptation.

enforcement constraint implies that when R's aggregate adaptation costs increase, the long-term value of the relationship between M and R must also increase in order to keep the relational contract within its self-enforcing range. In the next sections, we take this prediction to the data.

4. Data Description

4.1. Data

The data we use in this paper results from the combination of several data sets. We obtained airline ticket and flight information from the DB1B data, and ticket, market, and coupon data from RITA, both data sets from the Bureau of Transportation Statistics. These data contain not only information on the ticketing carriers, but also on the operating carriers and reporting carriers of each flight.⁴ We complement these data with information on aircraft type, operators, flight frequency and other route and flight characteristics (seats, number of flights, group of aircraft, distance flown, number of total passengers, and dummy of freighter flights), which we obtained from the T100-B41 and T100-B43 airline-aircraft data from the U.S. Department of Transportation. To merge the T100 and DB1B databases, we checked the identity of the ticketing, operating and reporting carrier of each flight.

We drop the freighter flights and the flights that have zero passengers from our data. We take the ticketing carrier identifier from DB1B market data because of the following

⁴ A ticketing carrier is the airline that sells airtickets to customers, whereas operating and reporting carriers are those operating the flight and reporting flight characteristics to the BTS.

two reasons: first, to identify and match with the operator from other data sets like DB1B ticket, coupon, and T100-B43, and second, in order to avoid overlooking code-sharing between airlines. In order to correctly identify contracts between major and regional airlines, we combine the merged DB1B and T100 datasets described above with the data from the Regional Airline Association (RAA), which provides the ownership type of each regional airline as well as the list of regional airlines, distinguished from charter airlines. We then merge this information with weather data on rainfall, snowfall and the number of freezing months per year-quarter (aggregated to the year-quarter level) from the National Oceanic & Atmospheric Administration.

By combining all these data sources, we obtain a rich data set that contains information at the major/regional/quarter and at the major/regional/route/quarter level, respectively. Following Forbes and Lederman (2009), we define a route as a set of one or more nonstop flights connecting the same two airports, irrespective of the flights' direction. We describe our two data sets below, and discuss our choice and use of variables in section 4.2.

The first data set, at the major/regional/quarter level, contains information on the number of joint routes, the number of flights, the number of seats, the highest value route served within the major-regional network, and weather conditions (average precipitation, snow, and number of freezing months across routes served by the dyad). All variables are computed for each quarter and year, and for each major/regional airline dyad.

The second data set contains information at the major/regional/route/quarter level. For each route, we create the following variables: a dummy for whether either airport in the route is a hub for the major airline in the dyad, a dummy for whether either airport in

the route is slot-controlled, the number of flights served in the route by the regional airline, the average value of a flight on the route (number of seats times the average price), the number of flights at the larger and smaller endpoints in the route, and weather conditions in the route—namely, the 1971-95 average snowfall, rain precipitation and number of freezing months from the National Oceanic and Atmospheric Administration (NOAA) evaluated at the endpoint airport in the route where they are highest (Forbes and Lederman 2009). We compute all of these variables for each quarter/year and for each major/regional/route triad.

4.2. Measures

The purpose of our empirical analysis is to assess whether following a negative shock, U.S. major airlines were more likely to drop or downsize routes outsourced to regional airlines whose *aggregate adaptation costs in the networks they operated for the major before the shock* were higher, as predicted by our relational adaptation model. As explained in the introduction, we focus on the exogenous shock represented by the financial crisis following the collapse of Lehman Brothers in September 2008. Accordingly, we define as our main dependent variable a dummy named “Survival - 9/2008” that takes value 1 if a given route operated by a regional airline on behalf of a major airline in 2006 (two years before the shock) is still operated by the same regional, on behalf of the same major, in 2010 (two years after the shock), and it takes value zero otherwise. Note that if a regional did not operate a given route in 2006, our survival variable excludes that route from the data.

Similarly, we create a second dummy variable, “Termination 9/2008,” that takes value 1 if the number of flights operated by a regional airline on behalf of a major airline

in a route has decreased between 2006 and 2010; and zero if the number of flights in the route has not decreased (that is, if it has stayed the same or increased). Again, routes that the regional airline did not fly in 2006 are left out of the sample. While the survival variable measures an extensive margin of adjustment in the major-regional relationships, the termination variable measures an intensive margin of adjustment.⁵

We provide graphical evidence on how the industry adjusted to the shock in Figure 1A and 1B below. Figure 1A shows that after the shock in 2008, the number of major-regional outsourcing relationships in the U.S. decreased, while the number of routes and flights ticketed by majors decreased slightly but far from the sudden and steep decrease experienced by the number of relationships between major and regional airlines. This evidence seems to indicate that the 2008 shock forced majors to reevaluate the amount of outsourcing relationships with existing regionals and restructure their network of regional flights accordingly.

<<Place Figure 1A here>>

<<Place Figure 1B here>>

Figure 1B focuses on the number of routes and flights outsourced to regionals and shows that, after the shock in 2008, even though the number of major-regional relationships in the U.S. decreased, the total number of routes and flights outsourced by major airlines to regional airlines increased. Thus, the evidence in Figure 1B indicates that while the 2008 shock did not decrease the use of outsourced regionals by major

⁵ See the Data Appendix for details on our treatment of airline mergers and exits during the 2006-2010 period.

airlines, it did push the majors to concentrate outsourced routes and flights into fewer regionals.

<<Place Table 1 here>>

Table 1 provides summary statistics for both the dependent and independent variables used in our study. We first note that our definition of the dependent variables, survival and termination, constrains the analysis to sample sizes of 6516 route-level observations. The probability that a route that was outsourced in 2006 survives in 2010 is 59.3%, whereas the probability that a route that was outsourced in 2006 has some of its flights terminated in 2010 is 62.2%.

Regarding the independent variables, we follow Forbes and Lederman (2009) and use adverse average weather conditions on a route — namely, snow (MAXsnowfall_r), precipitation (MAXprecipitation_r), and low number of freezing months (NFreezingmonths_r) as a proxy for the lack of clear skies — as exogenous proxies for the adaptation costs faced by the regional when bad weather hits that route. The underlying idea is that in routes characterized by more severe weather conditions, the major airline will more often require the regional to reschedule flights, thus inflating the regional's personnel and maintenance costs.

To measure aggregate weather conditions across a major-regional network, we computed the average of the key weather variables (maximum precipitation in inches, maximum snowfall in inches, and the number of freezing months in a route) across all routes flown by a regional for a major airline (AVEweatherSnow_ij, AVEweatherRain_ij, and AVEweatherFreez_ij, respectively).

As control variables, we use a number of route characteristics and major-regional network characteristics. In particular, we include variables that may drive outsourcing decisions regardless of network-level adaptation costs. Following Forbes and Lederman (2009), we include as route-level controls a dummy for whether either of the endpoints in a route is a hub for the major airline (*Dhubinroute_ir*), a dummy for whether either airport is slot-controlled (*slot_r*), and the total number of flights at the largest and smallest endpoints of the route (*flight_largeendpoint_ijr* and *flight_smallendpoint_ijr*). These variables may capture the extent to which a given route is embedded in the major's network. In turn, a route's embeddedness increases its strategic importance and the need for adaptation on the route, both of which may affect outsourcing for any given level of network-level adaptation costs. As additional route-level controls, we include the total number of flights operated by regional *j* in route *r* for major *i* (*NFlight_ijr*), and the average value of those flights (*AVEValue_ijr*), and the total number of flights at the largest and smallest endpoints of the route (*flight_largeendpoint_ijr* and *flight_smallendpoint_ijr*).

<<Place Table 2 here>>

<<Place Table 3 here>>

While Table 1 shows summary statistics for all of our independent variables, it does not provide any information on the characteristics of the different major-regional networks. For this reason, we complement Table 1 by providing information on the thickness and spread of the networks of major airlines in Table 2, and on differences in average network weather in Table 3. Table 2 tabulates the number of routes outsourced by each of the major airlines to each of the regional airlines in our data set. Note that the

number of regional partners, as well as the number of outsourced routes, varies across major airlines. The same pattern appears to be true from the regional perspective. While most regionals work for all majors, some regionals tend to concentrate their operations on one or two major airlines. Take for example the case of American Airlines (AA), which outsources routes to 19 different regionals but uses most intensively Envoy Air (MQ). Similarly, Envoy Air works mostly for American Airlines but also operates some routes for the other five major airlines in our data. Similar patterns characterize United Airlines (UA) and SkyWest Airlines (OO). See Figure 2 for an illustration of the networks of outsourced routes recently operated by SkyWest for several major airlines.

<<Place Figure 2 here>>

Given the large variation in network size across major airlines, we present in Table 3 summary statistics for the average network weather variables in our base year, 2006, for each major airline. Table 3 also reports the number of regional airlines with whom each major airline had a relationship during that year. The number of relationships in 2006 ranged between 16 (US Airways) and 19 (AA, Delta and United Airlines). Table 3 shows that there is a lot of variation in network weather variables (snow, rain, and number of freezing months) even within a major airline across its different regional networks. We exploit this variation later in our empirical analysis.

5. Empirical Methodology and Main Results

5.1 Empirical Methodology

Given our theoretical model from section 3, our hypothesis is that if a major and a regional airlines have entered a relational adaptation contract, the long-term value of their relationship must be higher the worse the weather conditions in all the regional routes they operate together, and therefore, the larger the cost of honoring the relational adaptation contract. Hence, we would expect major airlines to be more reluctant to drop or downsize routes belonging to a regional network with worse average weather, following the 2008 shock.

Therefore, our main empirical specification is a traditional linear probability model, estimated by OLS,⁶ such that:

$$Survival_{ijr} = \alpha + \beta AdaptationCost_{ij} + \gamma X_{ijr} + \delta_i + \mu_j + \theta_r + \varepsilon_{ijr},$$

$AdaptationCost_{ij}$ is the aggregate cost of adapting schedules across all routes jointly operated by major i and regional j in 2006, before the shock, and ε_{ijr} is a normally distributed and iid error term. X_{ijr} is a vector of observable characteristics of the ij relationship in route r in 2006, which includes route-level adaptation cost. Finally, δ_i , μ_j and θ_r are major airline, regional airline, and route fixed effects, respectively, which are

⁶ We choose to test our prediction with linear probability model and OLS because the number of fixed effects increases rapidly and so we want to avoid changes in methodology throughout the empirical results. Using probit for those specifications with no or few fixed effects does not qualitatively change our results.

included in the regression to control for unobservable components that are common across routes and common across airlines within a route. These fixed effects are crucial to our experimental design, because they allow us to compare the probability of continuation of two outsourced routes belonging to regional networks with different continuation values *for the same major airline and vice versa*.

Under the relational adaptation hypothesis, we expect $\beta > 0$. Routes belonging to networks with higher average adaptation costs are more likely to survive the Lehman Brothers shock because prior to the shock, those routes were allocated by major airlines to regionals whose relationship with the majors had a higher continuation value. Notice that under a “spot adaptation” hypothesis, we would expect instead $\beta = 0$, because absent relational contracts, outsourcing of a route should only depend on adaptation costs at the route level, as in Forbes and Lederman (2009), and not on network-level adaptation costs.

Because we do not directly observe $AdaptationCost_{ij}$, we use our measures of bad weather within the network ij as a proxy for $AdaptationCost_{ij}$ in the above equation. We rely on Forbes and Lederman (2009) to argue that there is a positive correlation between adaptation costs on a route and the incidence of bad weather on that route, so that bad weather aggregated at the network level is indeed a good proxy for network-level adaptation costs.

In this setting, we are consistently estimating β if the impact of the 2008 financial crisis on the survival probability of different routes in our sample is uncorrelated to route characteristics that determined the formation of major-regional networks prior to the

shock, in 2006. Formally, under the relational adaptation hypothesis and our specification above, our identification assumption is that $cov(\varepsilon_{ijr}, Weather_{ij}) = 0$.⁷

While there is no apparent reason to believe that routes in regional networks with worse weather were less likely to be affected by the financial crisis relative to routes in networks with better weather, we can think of two potential reasons why it may be that $cov(\varepsilon_{ijr}, Weather_{ij}) \neq 0$. The first reason is selection: a major may prefer to assign bad weather networks to regionals that do not renege on adaptation decisions, and a regional may prefer to operate a bad weather network for majors that do not renege on the relational adaptation bonus. The fact that networks are endogenously formed before the shock is not a problem per se—indeed, it is precisely this endogenous selection that enables us to use adaptation costs, proxied by bad weather, as a measure of the continuation value of major-regional relationships. However, selection may be a problem if it occurs in anticipation of the 2008 financial crisis shock. Our specification deals with this potential problem in two ways. First, we observe and code outsourcing relationships two years prior to the shock, so observed regional networks are unlikely to be formed in the anticipation of the 2008 financial crisis. Second, we use major and regional airline fixed effects, as well as combinations of major/route and regional/route fixed effects, in our robustness checks specifications.

The second potential reason why it may be that $cov(\varepsilon_{ijr}, Weather_{ij}) \neq 0$ is that our measures of network-level weather may be correlated with route-level weather (a route from a network with average bad weather is more likely to be itself a bad-weather route). In turn, a route with worse weather may be more likely to be cut after the shock because

⁷ $Weather_{ij}$ is defined as average weather in the major-regional network i and j .

it has higher adaptation costs and hence it is less profitable. To control for this second source of endogeneity, we include route-level weather, as well as route fixed effects, in our baseline regressions.

Thus, our identification assumption is that conditional on route and network characteristics, in the absence of a shock the profitability of routes should be rather unrelated to differences in adaptation costs across networks. After the 2008 negative shock, and if relational adaptation matters, we should find that major airlines are more likely to preserve outsourced routes that belong to regional networks with higher average adaptation costs, because those are the major-regional relationships with higher continuation value, so the major does not want to jeopardize them by reducing their size.

In section 5.2, we present our main results, which provide evidence on how average weather conditions across a major-regional network affected the survival and termination of routes in that network following the Lehman Brothers shock.

5.2. Main Results: Route survival as a function of Relational Adaptation

Tables 4 and 5 below report the effect of network weather conditions on the survival and termination of routes following the 2008 Lehman Brothers shock. Our independent variables are all divided by their own standard deviation in order to provide easier-to-interpret coefficients. We provide results without fixed effects in columns 1 and 4, with major airline fixed effects in columns 2 and 5, and with regional airline fixed effects in columns 3 and 6. In all specifications, standard errors are clustered at the major-regional dyad.

<<Place Tables 4 and 5 here>>

The results are consistent with our relational adaptation hypothesis. Routes in networks characterized on average by higher precipitation and more abundant snow, and by a lower number of freezing months and hence less clear skies, are more likely to survive (Table 4), and less likely to see their flights reduced (Table 5), following the Lehman Brothers shock. These results support our theoretical prediction that major airlines are more reluctant to restructure routes when their overall relationship with the regional airline serving those routes is valuable. The relationship's value is measured by the aggregate temptation to renege on a relational adaptation contract, as proxied by network-level adverse weather conditions, because for those contracts to be self-enforcing, higher-temptation outsourced networks must also have a higher long-term value.

Notice that the results are consistent regardless of the type of fixed effects included in the specification, indicating that our hypothesis is supported across major and regional airlines, within major airlines across their outsourced regional networks, and within regional airlines across the major airlines that contract with them.

Because our independent variables are standardized by their individual standard deviation, the interpretation of our empirical results is straightforward. Take, for example, column 2 in Tables 4 and 5, which includes major airline fixed effects. Our results show that a one standard deviation increase in the average precipitation across a major-regional network increases the probability of survival of a route by 16 percentage points, and decreases the probability of flight termination in a route by 8.5 percentage points.

When paying attention to the rest of independent variables included in the analysis to control for potential confounders, we find interesting and yet intuitive results. On the one

hand, we find that the number of flights in a route, and whether an airport in the route is slot controlled, have a statistically significant positive effect on route survival. On the other hand, our results show that a route's average value and the distance between its endpoints decrease its likelihood of survival. On one hand, routes with higher average value may be more likely to be vertically integrated after the shock (more below), and on the other hand, airline passengers may dislike longer flights in regional airlines due to the smaller aircrafts used. Finally, if anything, routes with a hub at an endpoint are less likely to survive or more likely to experience a reduction in the number of flights. This can be explained by the fact that, as shown by Forbes and Lederman (2009), routes with a hub at an endpoint may be more important in order for major carriers to achieve coordination with other flights and therefore, they may be more likely to be integrated after the 2008 shock.

<<Place Tables 6 and 7 here>>

Tables 6 and 7 reproduce the analysis in Tables 4 and 5, introducing route-level weather variables as independent variables, in addition to the network-level weather variables used previously. Our results show that route-level weather variables are statistically insignificant and that if anything, they have a mildly negative effect on survival, as one would expect given that high route-level adaptation costs reduce route profitability. In contrast, the network-level weather variables are still statistically significant and their signs are fully consistent with our relational adaptation hypothesis. Therefore we can conclude that our original results in Tables 4 and 5 are not due to correlations between route-level weather and network-level weather.

<<Place Tables 8 and 9 here>>

Tables 8 and 9 add route fixed effects to our specifications in Tables 4 and 5. Route fixed effects not only allow us to rule out differences in weather across routes as a potential explanation for our findings (as in Tables 6 and 7), but also allow us to control for any cross-route differences that may affect their post-shock survival. The regressions reported in Tables 8 and 9 hold the route constant, and exploit variation in survival and network-level weather across pairs of major and regional airlines that jointly operate that route. The results are entirely consistent with those in Tables 4-7. Therefore, it seems fair to conclude that the error term in our regressions (that is, route-specific random effects of the financial crisis) is uncorrelated with our main explanatory variable, that is, network-level average weather before the crisis.

<<Place Tables 10 and 11 here>>

We conclude this section by presenting Tables 10 and 11. These tables use the same specifications in our original Tables 4 and 5, with the difference that we now introduce major-route fixed effects (columns 1 and 3) and regional-route fixed effects (columns 2 and 4) for both our survival and termination dependent variables. The rationale behind the use of these new sets of fixed effects is to control for the possibility that the 2008 shock may have differently affected different airlines in different routes. Therefore, the specifications in Tables 10 and 11 only exploit variation within an airline-route dyad—that is, variation coming from majors that use more than one regional in a given route or regionals that operate the same route for more than one major.

The results in Tables 10 and 11 are again largely consistent with our findings in Tables 4 to 9. Routes outsourced to regional airlines are more likely to survive after the 2008 shock if they belong to networks that had worse weather (that is, higher precipitation and snow, and fewer freezing months per year) prior to the shock. If anything, the precipitation and snow coefficients in columns 2 and 4 of Table 11 are statistically less significant than those in previous table, although they remain large and negative.

6. Robustness checks and Additional Margins of Adjustments

6.1. Placebo Test: Survival of Outsourced Routes between 2003 and 2006.

Since the Lehman Brothers shock does not seem to have differentially affected different groups of routes, we cannot use a differences-in-differences estimation approach in our study. To test for whether the effects of network-level adaptation costs on route survival and size documented in tables 4-11 are really driven by the Lehman Brothers shock, we construct equivalent survival and termination variables taking as the initial and final years 2003 and 2006 — respectively, two years after the September 11 terrorist attacks and two years prior to Lehman Brothers September 2008 shock. Because no shock occurred between 2003 and 2006, we would expect our network-level weather variables, and hence the value of relational adaptation contracts, not to affect the probability of routes' survival and termination. Tables 12 and 13 below present our placebo test.

<<Place Tables 12 and 13 here>>

Table 12 provides evidence on route survival probability between 2003 and 2006. On the one hand, our results show that a route is more likely to survive when one of the airports in the route is a hub to the major airline, and when the number of flights operated by the major and the regional in that route, and the average value of flights in the route, increases. On the other hand, we find no statistically significant relationship between the average snow, precipitation and number of freezing months in the major-regional network and the probability of route survival. If anything, we find a weak negative relationship between precipitation levels and survival (columns 2 and 5), as well as a positive correlation between freezing months and route survival (columns 3 and 6).

Table 13 repeats the analysis in Table 12 with flight termination in the 2003 to 2006 period as our dependent variable. Our findings in Table 13 are consistent with those in Table 12 in that overall, average network weather conditions do not have a statistically significant effect on the likelihood that flights on a route are reduced.

Therefore, our placebo test suggests that absent an exogenous shock, there is no statistical correlation between the average weather conditions of a major-regional network and the survival and termination probabilities of a route within that same network. This evidence corroborates our hypothesis that the Lehman Brothers events unexpectedly forced major airlines to restructure their portfolios of regional routes, and are therefore appropriate “stress tests” for assessing the existence and significance of relational adaptation contracts.

6.2. The 9/11 Terrorist Attacks

As a robustness check, in this section we explore the relation between network-level adaptation costs and the survival of outsourced routes in the aftermath of the 9/11 terrorist attacks in 2001. It is important to notice that the U.S. airline industry in 2001 differed much from the U.S. airline industry in 2008, and also the nature and effects of the 9/11 was different from that of the Lehman Brothers shock. In particular, we notice that outsourcing to regional partners was less likely in 2001 than in 2008. Figure 1B shows that not only there were fewer major-regional relationships (albeit with a higher number of major airlines in operation) in 2001, but also, the number of outsourced routes and flights per regional in 2001 was far smaller than in 2008. Given that outsourced networks in 2001 were already thin, we do not expect major airlines to restructure them after the 9/11 shock as much as they did after the Lehman Brothers shock.

Tables 14 and 15 below display our results on survival and termination of routes before and after the 9/11 terrorist attacks. The dependent variables in these specifications are constructed in the same way as those in Tables 4 and 5. For each route that was outsourced to an independent regional in 1999 (two years prior to the 2001 shock), we compute dummies for whether the route was still outsourced to that regional (survival) and for whether its number of active flights had been reduced (termination) by 2003 (two years after the shock).

<<Place Tables 14 and 15 here>>

We find no robust statistical relationship between the average network weather variables and the survival and termination probabilities in the regressions without fixed

effects (columns 1 and 4). When using major airlines fixed effects in columns 2 and 5, we find that routes belonging to networks with more freezing months are more likely to survive and less likely to be terminated. When using regional airline fixed effects, we find that routes with more precipitation are less likely to survive and more likely to be terminated or reduced in size. We interpret these contradictory results as suggesting that the 9/11 shock had a weaker impact on the US airline industry than the 2008 shock at the core of our base analysis. This is consistent with our initial conjecture in that adjusting the major's networks is only possible through adjusting the number of outsourcing relationship if the major-regional networks are dense enough.

6.3. Alternative Margins of Adjustment

Our empirical analysis so far sheds light on the relationship between network-level adaptation costs and the termination or downsizing of outsourced routes in response to the 2008 financial crisis. Yet, there are other margins of post-shock network adjustment that have been so far left out of our study. When the outsourcing relationship between a major and a regional airline on a route or flight terminates, three things may happen. First, the major airline may stop flying the route altogether. Second, the major airline may fly the route with its own planes or through a vertically integrated regional company, as documented by Forbes and Lederman (2009). Third and finally, the major airline may continue to outsource the route but using a different regional airline, which may or may not have been previously used on that route.

In this last section, we investigate the relevance of the two latter margins of adjustment. Let us first study whether a route that was outsourced in 2006 is more likely to become integrated in 2010 as a result of the Lehman Brothers. Our results here add to the study of vertical integration by Forbes and Lederman (2009).

It is important to emphasize that the implications of our relational adaptation model for vertical integration are not as clear-cut as those for route survival and termination, because a major airline's potential cost savings from integration are dubious. On the one hand, labor costs tend to increase after integration because of the unionization of major airlines (Forbes and Lederman, 2009). On the other hand, major airlines may be forced to integrate some previously outsourced routes that were downsized after the shock, because no other independent regional could profitably operate those routes below a minimum number of flights. Thus, the net effect of the pre-shock expected value of major-regional relationship, as proxied by our weather variables, on post-shock vertical integration, appears to be an open empirical question.

To conduct this analysis, we use the same specifications used until now with two new dependent variables. First, we create a dummy variable, called *Integration*, that takes value 1 if, conditional on a route being operated by a regional airline in 2006, at least a flight in the route that was previously operated by the regional airline is operated by the major airline itself in 2010. We also create a second dependent variable, named *Integration2*, which results from conditioning *Integration* to at least one flight in the route being terminated after the shock (*Termination* =1). While the former variable checks whether any flight has been integrated, the latter restricts the analysis to those routes that experienced restructuring—that is, that saw the number of outsourced flights go down.

Tables 16 and 17 below report the effect of the average weather conditions in the major-regional network on the probability of integration after the 2008 shock.

<<Place tables 16 and 17 here>>

The results indicate that routes in networks with worse weather conditions, conditional on being outsourced in 2006, are less likely to become integrated and operated by the major airline after 2008. This is true within major airlines across outsourced regional airlines (columns 2 and 5 in Table 16), and within regional airlines across upstream major airlines (although less so due to conflicting signs between snowfall and number of freezing months). A one-standard deviation increase in snow decreases the probability of integration between 5% and 10.6%, and a one-standard deviation increase in precipitation decreases the probability of integration by roughly 6%. Results in Table 17 are very much consistent with those in Table 16 with Integration2 as dependent variable.

<<Place Table 18 here>>

Finally, we explore whether terminated outsourcing relationships are likely to be continued by another regional, and whether the continuing regional is an existing outsourcing partner or a new partner. To undertake this analysis, we divide routes per major airline and by the number of outsourcing regionals in 2006, and compute the probability of survival—that is, continuation of the route with the same airline that was operating it before the shock, as in our previous tables—the probability of continuation of the route with any regional partner that was not operating the route before the shock,—conditional on survival = 0, and the probability of continuation of the route through a

new partner (that is, a regional airline that was not operating any route for the major prior to the 2008 shock), conditional again on $\text{survival} = 0$. We provide the results of this exercise in Table 18.

We find that in routes where major airlines only used one regional airline before the shock, the likelihood of continuing operation of the route with another existing regional partner, conditional on termination of the previous partner on that route, is high. Even more importantly, the probability of using a new partner (that is, a partner not previously used in other routes) is zero across all regionals in our sample. When we look at the subsample of routes where major airlines used two regional airlines before the shock, we find again that the probability of continuing operations on a route with an existing regional partner after severing an outsourcing relationship on that route, is high, and that the probability of using a completely new regional partner is zero in all cases except for Continental, where it is positive but small (2.22%). When we examine cases where major airlines use more than two regionals before the shock, we find that the likelihood of continued operation of a route is very high (close to 100% for all major airlines except for Northwest⁸) but again, continuation under a new outsourcing partner (not used before on any route) is zero for all major airlines except for US Airways.

In conclusion, the beginning of the financial crisis in 2008 marked by the disappearance of Lehman Brothers forced US airlines to redesign their networks of outsourcing relationships and we found that they did so in the three following ways. First, they terminated their existing outsourcing agreements on routes that were outsourced to

⁸ This is mainly driven by the merger of Delta and Northwest, and how Delta may have used their own partner to continue operations. See in the Data Appendix how we treat in our data the Delta-Northwest merger for clarification.

regionals with low continuation value, proxied in our analysis by the network-level average weather. Second, they integrated routes that were previously outsourced to regionals with low continuation value. Third and last, the majors reallocated most of the terminated routes to pre-existing partners that had higher continuation value. All these three margins of adjustment confirm the importance of the value of relational contracting for efficient adaptation in the US airline industry between major and regional airlines.

7. Conclusion

In this paper we study the value of relational adaptation in outsourcing relationships, using data from the US airline industry. Our theoretical model shows that for relational adaptation contracts between major and regional airlines to be self-enforcing, the long-term value of the relationship must be at least as large as the regional's cost of adapting flight schedules across joint routes. Thus, when facing a shock that forces it to terminate some routes, the major is more likely to preserve routes outsourced to regional airlines that have higher adaptation costs, as the value of the major's relationship with those regionals is larger.

In our empirical analysis, we analyze the evolution of major-regional airlines networks in the U.S. around the 2008 financial crisis, and we find that consistent with our theoretical predictions, regional routes belonging to networks with worse average weather, and hence higher adaptation costs, were more likely to survive after the shock. This finding is robust to the inclusion of route-level weather variables as well as route and airline fixed effects.

While it is often argued that both adaptation to unforeseen contingencies and informal, self-enforcing agreements are of fundamental importance to the success of inter-firm collaborations, there is still little evidence supporting these claims. Our hope is that the evidence and methodology provided by this study will contribute to shed light on these important phenomena. We also hope that our work will inspire future research that may further expand our understanding of how relational contracts help solving adaptation problems that spot and formal contracts fail to address.

References

- Arruñada, B., L. Garicano, and L. Vázquez, (2001), “Contractual Allocation of Decision Rights and Incentives: The Case of Automobile Distribution,” *Journal of Law, Economics and Organization* 17: 257-284.
- Bajari, P., and S. Tadelis, (2001), “Incentives versus Transaction Costs: A Theory of Procurement Contracts,” *The RAND Journal of Economics*, Vol. 32, No. 3 (Autumn, 2001), pp. 387-407.
- Baker, G., R. Gibbons, and K. J. Murphy, (2011), “Relational Adaptation” MIT working paper.
- Barron, D., R. Gibbons, R. Gil, and K. J. Murphy, (2015), “Relational Adaptation Under Reel Authority,” mimeograph.
- Chakravarty, S., and B. MacLeod, (2009), “Contracting in the Shadow of the Law,” *RAND Journal of Economics* 40: 533-557.

- Crocker, K., and K. Reynolds, (1993), "The Efficiency of Incomplete Contracts: An Empirical Analysis of Air Force Engine Procurement," *RAND Journal of Economics* 24: 126-146.
- Fast, N., and N. Berg, (1975), "The Lincoln Electric Company," Harvard Business School Case #376-028.
- Forbes, S., and M. Lederman, (2009), "Adaptation and Vertical Integration in the Airline Industry," *American Economic Review* 99: 1831-1849.
- Forbes, S., and M. Lederman, (2010), "Does Vertical Integration Affect Firm Performance? Evidence from the Airline Industry," *RAND Journal of Economics* 41: 765-790.
- Foss, N, (2003), "Selective Intervention and Internal Hybrids: Interpreting and Learning from the Rise and Decline of the Oticon Spaghetti Organization," *Organization Science* 14: 331-349.
- Gil, R., and G. Zanarone, (2015a), "On the Determinants and Consequences of Informal Contracting," Johns Hopkins Carey Business School working paper.
- Gil, R., and G. Zanarone, (2015b), "New Frontiers in the Empirical Research on Informal Contracting," *Journal of Institutional and Theoretical Economics* (forthcoming).
- Hart, O., and J. Moore, (1988), "Incomplete Contracts and Renegotiation," *Econometrica* 56: 755-85.
- Helper, S., and R. Henderson, (2014), "Management Practices, Relational Contracts, and the Decline of General Motors," *Journal of Economic Perspectives* 28: 49-72.

- Levin, J., (2002), "Multilateral Contracting and the Employment Relationship," *Quarterly Journal of Economics* 117: 1075-1103.
- Levin, J., (2003), "Relational Incentive Contracts," *American Economic Review* 93: 835-57.
- Macaulay, S., (1963), "Non-Contractual Relations in Business: A Preliminary Study," *American Sociological Review* 28: 1-19.
- Macchiavello, R., and A. Morjaria, (2015), "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," forthcoming in *American Economic Review*.
- MacLeod, B., (2007), "Reputations, Relationships, and Contract Enforcement," *Journal of Economic Literature* 45: 595-628.
- MacLeod, B., and J. Malcomson, (1989), "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica* 57: 447-480.
- Malcomson, J., (2013), "Relational Incentive Contracts," In Gibbons, R., and J. Roberts (editors), *The Handbook of Organizational Economics*, Princeton: Princeton University Press.
- Masten, S., and K. Crocker, (1985), "Efficient Adaptation in Long-Term Contracts: Take-or-Pay Provisions for Natural Gas," *American Economic Review* 75: 1083-1093.
- Schwartz, A., (1992) "Relational Contracts in the Courts: An Analysis of Incomplete Agreements and Judicial Strategies," *Journal of Legal Studies* 21: 271-318.
- Schummer, J., and Vohray, R., (2013), "Assignment of Arrival Slots," *American Economic Journal: Microeconomics*, 5 (2): 164-185.

- Simon, H., (1951), "A Formal Theory of the Employment Relationship," *Econometrica*, Vol. 19, No. 3 (Jul., 1951), pp. 293-305.
- Vossen, T., and, Ball, M. (2006), "Optimization and mediated bartering models for ground delay programs," *Naval Research Logistics* Vol. 53, Issue 1, pages 75–90.
- Williamson, O., (1975), "Markets and Hierarchies: Analysis and Antitrust Implications" New York, NY: Free Press.
- Williamson, O., (1991), "Comparative Economic Organization: The Analysis of Discrete Structural Alternatives," *Administrative Science Quarterly* 36: 269-96.
- Zanarone, G. (2009), "Vertical Restraints and the Law: Evidence from Automobile Franchising," *Journal of Law and Economics* 52: 691-700.
- Zanarone, G., (2013), "Contract Adaptation under Legal Constraints," *Journal of Law, Economics and Organization* 29: 799-834.

Data Appendix

We made a few assumptions in order to assemble our data when facing instances of missing information and the various mergers and exits that occurred in the US airline industry between 1999 and 2010.

To create our main dependent variables, Survival and Termination between 2006 and 2010, as well as the dependent variables in our placebo test (Survival and Termination between 2003 and 2006) and in the examination of the 2001 terrorist attacks (Survival and Termination between 1999 and 2003), we utilize the information available in the DB1B data set to the best of our ability. In particular, we code our major-regional airline interaction based on the ticketing carrier code (major airline) and the operating carrier/reporting carrier code (actual operator). Because some (about 20%) observations in the DB1B data matched with T100 do not have an operating carrier number (code "99") but do have a reporting carrier number, we replace the operating carrier with the reporting carrier for the observations with a missing operating carrier. According to BTS,⁹ the reporting carrier is usually the operating carrier of the first segment of an itinerary. Because we only use nonstop flights and the segment for a nonstop flight is from the origin airport where you take off to the destination airport, this assumption should not suppose a problem.

So far as our main analysis (2008 Lehman Brothers shock) is concerned, these are the mergers and exits that we encountered:

- (1) Delta (DL) and Northwest (NW) merged in 2008 and were operating only under DL in 2010. We assume that a route outsourced by NW to a given regional in 2006 survived to 2010 if we observe DL outsourcing that route to the same regional in 2010.
- (2) Republic AL (RW) and Midwest AL (YX) merged. Even though Republic AL survived, it changed its airline code to Midwest AL (YX). We apply same assumption as for DL and NW merger.
- (3) United (UA) and Continental (CO) announced their merger in 2010 but they were not able to close it until 2012. Hence, this merger does not affect our data and empirical analysis.
- (4) Pinnacle AL and Colgan AL merged in 2008 but operated separately through 2010.

⁹

https://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/subject_areas/airline_information/accounting_and_reporting_directives/number_224.html

- (5) Pinnacle AL and Mesaba AL merged in 2008 but operated separately through 2010.
- (6) Skywest, AS AL and ExpressJet AL merged in 2008 but operated separately through 2010.
- (7) A number of regional airlines declared bankruptcy but continued operating afterwards. These are Sky Airlines in 2008, Mesa Airline in 2010, Skybus Airline in 2008, Arrow Air in 2010, Sun country Air in 2008, AirMidwest in 2008, and Big sky in 2008.
- (8) Two regional airlines ceased operations and exited in 2008: AirMidwest (used by US Airways), and Big Sky (used by Delta). Their relationships with major US Airways and Delta appeared as not surviving in our sample.

So far as our placebo test is concerned (period between 2003 and 2006), here are the operations we have identified and the corresponding assumptions we have made in assembling our our data:

- (1) Two regionals named Republic AL (RW) and Shuttle America (S5) merged in 2005 into Republic AL. S5 appeared only once in 2003. We classified a route that was outsourced by a given major to S5 in 2003 and to RW in 2006 as a route outsourced to S5 in 2006.
- (2) Skywest and Atlantic Southeast Airline merged in 2005. Despite that, both operated separately in 2006.
- (3) US Airways' acquired America West Airlines (HP) in 2005. HP ceased operations in 2005, but HP still appears as HP (not US Airways) in 2006.

So far as our investigation of the terrorist attacks of 9/11, we only encountered the merger between American Airlines (AA) and Trans World Airlines (TWA) in 2001. We assume that a route outsourced by TWA to a given regional in 1999 survived to 2003 if we observe AA outsourcing that route to the same regional in 2003.