BARRIERS TO ENTRY IN THE AIRLINE INDUSTRY: A MULTIDIMENSIONAL REGRESSION-DISCONTINUITY ANALYSIS OF AIR-21

Connan Snider and Jonathan W. Williams*

Abstract—We investigate the success of legislation aimed at increasing competition at highly concentrated U.S. airports, mainly by forcing these airports to increase the availability of scarce facilities. We use a multidimensional regression-discontinuity approach to exploit a sharp discontinuity in the law's implementation and identify its effects. We find that fares decrease by 13.4% (20.2%) in markets with one (both) end point(s) covered. Approximately half of the decline is driven by the entry of low-cost carriers. We find little evidence that the fare declines were accompanied by a diminished quality of service, and passenger volumes increased, which suggests the legislation improved consumer welfare.

I. Introduction

ONE of the most enduring features of the postderegulation U.S. airline industry has been the hub premium—the premium over average fares paid by passengers traveling to and from airports dominated by a single large hub carrier. Though this phenomenon has been widely documented (Borenstein, 1989), its causes and consequences are still in question. To the extent that higher fares result from the exercise of market power, they are detrimental to consumer welfare and efficiency. However, there is substantial evidence that consumers value the large route network and high frequencies that dominant carriers often provide (Berry, 1990). To the extent that high prices derive from these quality factors, they benefit consumers. Airport facilities, a necessary input for the provision of air service, are increasingly scarce due in large part to government regulation handicapping their' ability to finance facilities expansion in response to increasing demand. As a consequence of this increased scarcity, airport-level concentration and market power go hand-in-hand with the airline scope and scale that consumers value and drive down costs. Given this, the relative contribution, and the optimality of the balance, of these factors is an empirical question.

In 2000, the U.S. Congress enacted the Wendell H. Ford Aviation Investment and Reform Act for the 21st Century (AIR-21). A primary directive of the bill was to require airports, above a given level of concentration, to take concrete steps to ensure that new entrants had ample access to airport facilities.[1] Airport compliance requires filing a Competition Plan with the Federal Aviation Administration (FAA), detailing the steps taken. The Federal Aeronautics Administration then reviews the plan and releases federal funding contingent on a satisfactory plan.

In this paper we empirically evaluate the impact of AIR-21 on prices, passenger volumes, quality, and market structure to investigate the importance of access to airport facilities as barriers to entry in the airline industry. The nature of the implementation of AIR-21 is useful for solving identification problems common in industrial organization studies of competition and market structure and present in our context. The problem is that elements of market structure (e.g., concentration, low-cost carrier presence) are determined simultaneously with the level of competition and usually depend on common, market-specific unobserved factors (e.g., demand elasticities or network economies associated with airport geography). We use the design of AIR-21 to formulate a differences in differences (diff-in-diff) and regression discontinuity (RD) solution to these problems.

We first argue that the AIR-21 mandates were enforced and effectively reduced barriers to entry at covered airports. This generates rarely available variation, with a plausibly known direction, over time in barriers to entry within markets. This allows us to control for time-invariant, market-specific factors using standard panel techniques. Second, having contemporaneous treatment and control groups allows us to use diff-in-diff to address aggregate and market-specific variation in these factors over time.

There are still likely to be selection problems associated with using the full sample for identification. Berry and Jia (2010), observing lower fares and diminished carrier profit margins between the end of the 1990s and the mid-2000s, estimate discrete choice demand systems separately for 1997 and 2005 and conclude that increased passenger price sensitivity, combined with increased penetration of low-cost carriers, was responsible for the change. Since airport concentration itself is likely highly correlated with product quality, the time-varying relative valuations of quality found by Berry and Jia (2010) likely interact with our determinant of treatment. This causes differing average trends for treated and untreated markets, invalidating the simple diff-in-diff approach. We solve this and any similar such problems by arguing that while there is likely a selection problem associated with highly concentrated airports, there is no such problem locally around the 50% two-carrier concentration level specified by AIR-21. This allows us to develop an RD estimator for the local average treatment effects associated with AIR-21. Essentially

Received for publication March 25, 2013. Revision accepted for publication February 13, 2014. Editor: Philippe Aghion.

* Snider: UCLA; Williams: University of North Carolina.

This paper has benefited from discussions with Severin Borenstein, Jan Brueckner, Dennis Carlton, Federico Ciliberto, Alon Eizenburg, Silke Forbes, Alfred Galichon, Mara Lederman, Yehua Li, Nancy Rose, Nick Rupp, Stephen Ryan, and Steven Stern. Seminar participants at Clemson University, London Business School, NYU, the Sauder School at the University of British Columbia, and UCLA, as well as conference participants at the 2010 IIOC and 2011 NBER IO Summer Institute, provided valuable comments. We are grateful to Sharon Glasgow and Andrea Toney of the FAA, as well as Liying Gu, Deborah McElroy, and A. J. Muldoon of the ACI-NA for assistance in collecting the data for this project and providing valuable industry insight. All remaining errors are our own.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00455.

[1] The law applied to airports in which the top two airlines accounted for over 50% of total enplanements at the airport.

The Review of Economics and Statistics, December 2015, 97(5): 1002–1022

doi:10.1162/REST_a_00455

we assume that the distribution of unobservables for a randomly selected market just below the cutoff is identical to a randomly selected market just above the cutoff.

The design of AIR-21 also helps us dismiss concerns about manipulation of the predictors of coverage. Airport coverage is determined by traffic data from two years prior to coverage, making coverage dependent on the past actions of the carriers, which are not subject to manipulation. Given the complexity of airline pricing decisions, it also seems unlikely that carriers would adjust fare-setting behavior to manipulate enplanements at the airport level. Nevertheless, we design an informal test of manipulation. The test is based on the observation that for a given two-firm airport concentration level, airports with higher one-firm concentration levels would be more likely to see manipulation, since a single carrier has more control over the coverage variable. This test shows no evidence of manipulation.

Ultimately we implement two RD estimators. The first takes the boundary dummy approach of Black (1999), using observations from progressively smaller windows around the treatment cutoffs. The second is a novel, true RD estimator. Since airline markets necessarily involve both an origin and a destination airport, we naturally have to consider two predictor variables and four treatment/control groups when defining our treatment effects; both the origin and destination are treated, just the origin is treated, just the destination is treated, and neither is treated. Our approach is essentially to estimate the surface for each of these "quadrants" and look for breaks at the boundaries of those quadrants. Our implementation of the two-dimensional RD estimator is novel in that it allows substantial heterogeneity by permitting the average treatment effect of one end point to vary, essentially, according to the identity of the other end-point airport.

Our relatively clean identification strategy represents a contribution to the literature on airline market structure and the importance of barriers to entry more generally. A typical structural study of entry and market structure in concentrated industries (Bresnahan & Reiss, 1991; Berry, 1992; Mazzeo, 2002; Seim, 2006; Ciliberto & Tamer, 2009) looks at firm choices and uses a revealed preference approach to infer entry barriers. This approach requires the economist to rely on many restrictions of the empirical model derived from economic theory. In the end, barriers to entry are often measured as the residual that justifies the cross-market heterogeneity in market structures. Our approach instead uses a known source of exogenous variation in entry barriers to investigate their effects on market outcomes and requires little necessary theoretical structure. The minimal structural requirements are useful for an industry as complex as the airline industry, where it is difficult to construct a model that is simultaneously rich enough to be quantitatively relevant and tractable enough to be amenable to analysis. Also, our focus on outcomes makes our results directly relevant for policy.

To preview results, we find AIR-21 had a substantial, and evidently positive, impact on competition and fares in the airline industry. We find that markets for which one of the end-point airports was subject to AIR-21 have seen price declines of 13% on average. Markets for which both end points were subject to the mandates have seen price declines of around 20% on average. These price declines were associated with no economically and statistically significant changes in measures of quality, with one exception. We find that the on-time performance of carriers at covered airports decreased. This is not particularly surprising, as we identify increased low-cost carrier penetration as a driving force behind the declines in fares, reinforcing the findings of Mayer and Sinai (2003) that higher airport concentration leads to better on-time performance as large carriers, presumably, internalize more of the congestion costs. We also find that the magnitude of the decline in fares is greater for carriers with a large presence at an airport than for other carriers. This suggests that AIR-21 was successful at reducing the hub premiums that Borenstein (1989) identified. While the delay results suggest some quality degradation on some dimensions, the magnitudes of the price declines and accompanying increases in passenger volumes indicate the legislation was, overall, welfare enhancing.

The implications of our findings are not limited to retrospective policy evaluation. The competition provisions of AIR-21 address the allocation of a given stock of airport facilities. A more natural policy prescription might be to encourage the expansion of this stock. However, since 2001, passenger facility charges (PFCs), the per boarding passenger tax levied by the airport and a primary source of funding for airport investment projects for large airports and the primary backing for the external finance of such projects, have been federally capped at $4.50, in spite of protest by airport authorities. In a letter to President Obama dated August 18, 2011, Greg Pricipatio, president of the trade association Airports Council International–North America (ACI-NA), urged the president to lift the cap as part of his jobs creation plan, saying, "The $2 billion [PFC] raises a year is already committed to paying for recently completed projects or those currently underway." Resistance to PFC increases has come from large carriers that wish to avoid their share of the tax burden and simultaneously prevent expansion of their home airports and the accompanying increase in competition. The influence of these carriers is not surprising given their importance not only to the airport where they operate large hubs, but also to the communities, where they employ thousands of workers and provide a substantial tax base. While our results do not speak directly to the effect of airport expansion, they suggest this resistance may come at the cost of diminished competition and sharply higher fares for air travelers in those communities and may warrant more careful study.

The remainder of the paper is organized as follows. In section II, we provide some background on the airline industry and discuss AIR-21 in detail. The data are described in section III, and we document some basic patterns in the data over the policy period. In section IV, we discuss our identification strategy and the results of our analysis. Section V concludes and discusses possible extensions of our research.

II. The Aviation Investment and Reform Act for the Twenty-First Century

The Government Accounting Office (GAO) and Transportation Research Board (TRB) released a series of reports (GAO, 1989, 1990, 2001; TRB, 1999), bringing attention to the limited amount of competition at many major U.S. airports. These reports identified two types of barriers to entry in the airline industry that can limit competition and result in higher fares: operating and marketing.

Marketing barriers include loyalty programs intended to tie consumers to an airline, frequent flier programs, corporate incentive agreements, and travel agent commission overrides. A lack of data has limited the study of these type of barriers, with Lederman (2007, 2008) and Goolsbee and Syverson (2008) as notable exceptions. Lederman (2007, 2008) finds evidence that improvements in loyalty programs enhance demand and can explain a modest portion of the hub premium. Goolsbee and Syverson (2008) show that national carriers respond to the "threat of entry" by Southwest Airlines, a low-cost carrier, by lowering fares to strengthen consumer loyalties prior to entry of Southwest.

Operating barriers include limited access to boarding gates, ticket counters, baggage handling and storage facilities, and take-off and landing slots. Ciliberto and Williams (2010) were the first to directly link these operating barriers to the hub premium. Using unique data on carrier-specific access to boarding gates, they show that long-term exclusive-use leasing agreements for boarding gates limit competition and are a major driver of the hub premium. In this paper, we employ a unique identification strategy to examine the success of AIR-21 in reducing these operating barriers and encouraging competition at major U.S. airports. In the sections that follow, we discuss the details of AIR-21's design and implementation.

A. *Legislation and Airport Coverage*

In response to governmental, public, and academic concern with institutional barriers to entry in the airline industry, President Clinton signed into law AIR-21 on April 5, 2000. Section 155 of AIR-21 begins:

> The Congress makes the following findings:
>
> (1) Major airports must be available on a reasonable basis to all air carriers wishing to serve those airports.
>
> (2) 15 large hub airports today are each dominated by one air carrier, with each such carrier controlling more than 50 percent of the traffic at the hub.
>
> (3) The General Accounting Office has found that such levels of concentration lead to higher air fares.
>
> (4) The United States Government must take every step necessary to reduce those levels of concentration.
>
> (5) Consistent with air safety, spending at these airports must be directed at providing opportunities for carriers wishing to serve such facilities on a commercially viable basis.

Together findings 1, 4, and 5 demonstrate Congress's clear intentions to reduce concentration by encouraging additional entry at concentrated airports. To encourage airports' cooperation in opening up airports to "all air carriers wishing to serve those airports," Congress made federal sources of funding contingent on compliance:

> Beginning in fiscal year 2001, no passenger facility fee may be approved for a covered airport under section 40117 and no grant may be made under this subchapter for a covered airport unless the airport has submitted to the secretary a written competition plan in accordance with this subsection.

Passenger facility fees (commonly called PFCs) and airport improvement program (AIP) grants are the primary sources of federal funding for the industry and make up a significant portion of capital (including maintenance) budgets for major airports.[2] PFCs were first authorized by Congress in 1990 and are tied to projects to preserve and enhance safety, reduce noise pollution, and provide opportunities for greater competition between carriers. The PFC ceiling, the maximum fee allowed by law, was increased from \$1.00 to \$4.50 between 1990 and 2001. This ceiling has not been increased since AIR-21 and is not indexed for inflation. AIP grants are part of a federal program to help cover costs for approved capital projects aimed at increasing safety and capacity, as well as reducing environmental concerns.

A 2009 Airport Council International North America (ACI-NA) study found that over 40% of airports' capital funding is drawn from PFCs (21.7%) and AIP grants (22.2%).[3] PFCs alone have funded \$50 billion worth of airport capital investments since 1990, including the addition and maintenance of passenger boarding gates and runways necessary to accommodate additional entry. An additional 30% of airports' revenues come from bonds, which are often backed with future PFCs revenues. This substantial and stable revenue base allows airports to significantly lower the cost of borrowing and enjoy investment-grade ratings. While the quasi-public status of many airports makes it difficult to know their exact objectives, the strong dependence of airports' revenues on the federal government's control over the right to charge PFCs and distribute AIP grant funding seems to imply strong incentives for compliance. All airports covered by AIR-21 are required to file a competition plan with the FAA and the DOT, which must certify the plan as acceptable for funding to be released.[4]

[2] PFCs are charged by airlines at the time a ticket is purchased and are then transferred directly to the appropriate airports.

[3] A copy of the presentation describing this report is available from the authors on request.

[4] The 44 airports required by AIR-21 to file a competition plan are: ABQ, ANC, ATL, AUS, BNA, BUR, BWI, CLE, CLT, CVG, DAL, DCA, DEN, DFW, DTW, EWR, HOU, IAD, IAH, JAX, LAS, MDW, MEM, MIA, MKE, MSP, OAK, OGG, ONT, ORD, PBI, PHL, PHX, PIT, PVD, RNO, SAT, SDF, SFO, SJC, SJU, SLC, SMF, and STL. All but one of the airports were immediately covered by the retroactive nature of the legislation. The only airport to be covered later was LAS in 2005.

Congress also made it clear that competition “plans” were to be implemented:

> The Secretary shall review any plan submitted ... to ensure that it meets the requirements of this section, and shall review its implementation from time-to-time to ensure that each covered airport successfully implements its plan. ... The Secretary shall ensure that gates and other facilities are made available at costs that are fair and reasonable to air carriers at covered airports ... where a “majority-in-interest clause” of a contract or other agreement or arrangement inhibits the ability of the local airport authority to provide or build new gates or other facilities.

In conversations with those at the FAA assigned to approve and ensure implementation of the competition plans, we learned that approval was not a certainty for any plan. In many cases, the plans were significantly revised after discussions between the FAA, DOT, and airport authorities to ensure the plans meet the goals of the legislation. After filing the initial competition plan, airports were required to complete two updates (approximately eighteen months apart) that demonstrate significant progress toward implementing the competition plan. There are no mandatory steps after the second update for covered airports unless the airport denies a carrier access to airport facilities or significantly amends an existing leasing agreement or enacts a new master-leasing agreement.

Section 155 continues:

> A competition plan under this subsection shall include information on the availability of airport gates and related facilities, leasing and sub-leasing arrangements, gate-use requirements, patterns of air service, gate-assignment policy, financial constraints, airport controls over air- and ground-side capacity, whether the airport intends to build or acquire gates that would be used as common facilities, and airfare levels (as compiled by the Department of Transportation) compared to other large airports.

The typical competition plan ranges in length from 75 to 100 pages and contains a vast amount of information about the airport’s operations. Ciliberto and Williams (2010) use this information to demonstrate that Congress’s focus on equal access to sunk airport facilities is not completely misguided. Using cross-sectional variation in gate allocations and leasing terms, they can explain an economically significant fraction of the hub premium, with this fraction being larger at congested airports. In this paper, we focus on measuring any reduction in fares and the hub premium resulting from coverage of an airport by AIR-21.

To identify the impact of AIR-21 on the hub premium, and fares more generally, we exploit the sharp discontinuity in the relationship between coverage and concentration: “‘covered airport’ means a commercial service airport . . . that has more than .25 percent of the total number of passenger boardings each year at all such airports . . . at which one or two air carriers control more than 50 percent of the passenger boardings.”

These concentration thresholds create treatment and control groups, airports “very near” either side of the discontinuity, which can be used to measure the impact of the legislation on competition.[5] An airport is covered by the legislation if it qualifies in both the size and concentration dimensions.[6] In section IV, we discuss how we exploit this feature of the legislation using a regression discontinuity approach to measure a (local) treatment effect or impact from coverage at the concentration cutoff. Tables 1 and 2 show the two-firm enplanement concentration and the fraction of total domestic enplanement at covered and noncovered airports, respectively. While concentration and size are positively correlated, it is far from a perfect relationship. For example, Newark (EWR) is covered, while New York (JFK) is not. Similarly, San Francisco (SFO) is covered, while Los Angeles (LAX) is not.

B. *Implementation of Competition Plans*

Most of the competition plans and subsequent updates are available on each airport’s website. The details of each competition plan are too vast to review here. However, a 2006 FAA report highlights specific actions taken by airports in a variety of areas to increase competition.[7]

In terms of improving availability of gates and related facilities, airport responses included asserting control over underused gates, designating competition access committees, adopting more entry-friendly leasing terms, removing specific access protections for signatory carriers, and streamlining a forced accommodation process. There were a number of specific actions. Hartsfield-Jackson Atlanta International Airport (ATL), for example, invoked recapture authority to convert a leased gate to common use, Cincinnati-Northern Kentucky International Airport (CVG) negotiated conversion of exclusively leased gates to common and preferentially leased gates, and San Francisco International Airport (SFO) invoked a forced accommodation clause to ensure that temporary needs of new entrant airlines were met. In terms of subleasing agreements, covered airports also began to assert more control and oversight over sublease fees, terms, and conditions; impose sublease caps on administrative fees; review or preapprove subleases; and notify carriers of gates available for subleases.

Improving access to passenger boarding gates was clearly the focus of a large proportion of each competition plan.

[5] As with any other analysis examining treatment effects, the treatment must be exogenously applied. In the context of our study, endogeneity of treatment might arise if airports can lower the concentration of enplanements or total enplanements to avoid being covered by the legislation. In section IVB, we show that there is little or no support for the claim that enplanements were strategically manipulated by carriers to avoid coverage.

[6] The discontinuity along the size dimension also presents an opportunity to identify an effect from coverage, but the small number of airports near this cutoff limits our ability to exploit this feature of the law.

[7] This report is available through the FAA website: http://www.faa.gov/airports/aip/guidance_letters/media.

TABLE 1.—ENPLANEMENTS AND GATES FOR COVERED AIRPORTS

		Enplanements			Gates					
		United States %	Top 2%		Common %		Legacy %		LCC %	
Airport	Year Covered	Mean	Mean	Maximum	2001	2008	2001	2008	2001	2008
ABQ	2000	0.45%	61.24%	63.97%	26.09%	31.82%	39.13%	63.64%	34.78%	4.55%
ANC	2000	0.36%	55.23%	61.74%						
ATL	2000	5.85%	79.17%	82.18%	14.59%	15.08%	72.43%	73.37%	12.97%	11.56%
AUS	2000	0.51%	60.32%	61.80%	28.00%	16.00%	44.00%	52.00%	28.00%	32.00%
BNA	2000	0.64%	59.02%	63.03%	11.48%	9.84%	44.26%	44.26%	44.26%	45.90%
BUR	2000	0.36%	77.98%	83.54%	21.43%	7.14%	28.57%	35.71%	50.00%	57.14%
BWI	2000	1.42%	56.59%	65.95%						
CLE	2000	0.84%	58.97%	61.29%						
CLT	2000	1.82%	81.43%	86.84%	44.71%	48.35%	55.29%	51.65%	0.00%	0.00%
CVG	2000	1.52%	87.47%	92.87%						
DAL	2000	0.48%	97.79%	99.82%	18.75%	0.00%	15.63%	25.00%	65.63%	75.00%
DCA	2001	1.09%	44.06%	50.10%						
DEN	2000	2.82%	66.04%	72.44%						
DFW	2000	4.06%	77.14%	85.12%	5.47%	17.42%	89.06%	80.00%	5.47%	2.58%
DTW	2000	2.47%	72.71%	76.32%	5.47%	5.08%	84.38%	88.14%	10.16%	6.78%
EWR	2000	2.41%	59.77%	69.90%						
HOU	2000	0.61%	89.23%	92.19%						
IAD	2001	1.42%	53.70%	59.91%						
IAH	2000	2.50%	80.91%	86.12%						
JAX	2000	0.37%	46.52%	50.19%						
LAS	2005	2.68%	47.81%	52.40%						
MDW	2000	1.11%	77.90%	90.37%						
MEM	2000	0.80%	72.17%	77.10%						
MIA	2001	2.29%	57.22%	68.95%	21.65%	32.04%	74.23%	64.08%	4.12%	3.88%
MKE	2001	0.47%	49.87%	56.49%	17.95%	0.00%	48.72%	21.28%	33.33%	78.72%
MSP	2000	2.44%	75.89%	78.75%	9.52%	8.66%	86.90%	90.55%	3.57%	0.79%
OAK	2000	0.88%	72.26%	78.52%	12.50%	37.93%	16.67%	3.45%	70.83%	58.62%
OGG	2000	0.42%	60.00%	68.59%						
ONT	2000	0.48%	59.52%	61.44%						
ORD	2000	5.01%	67.79%	74.12%						
PBI	2000	0.46%	52.29%	58.64%	50.00%	53.13%	39.29%	34.38%	10.71%	12.50%
PHL	2000	1.92%	60.61%	65.66%						
PHX	2000	2.72%	66.85%	68.95%						
PIT	2000	1.22%	66.85%	81.65%						
PVD	2000	0.39%	56.71%	63.55%						
RNO	2000	0.38%	58.91%	62.61%						
SAT	2001	0.48%	57.94%	100.00%	16.67%	17.39%	54.17%	52.17%	29.17%	30.43%
SDF	2000	0.27%	45.50%	51.64%						
SFO	2000	2.44%	53.42%	56.29%	36.14%	37.04%	60.24%	59.26%	3.61%	3.70%
SJC	2000	0.82%	57.53%	64.18%						
SJU	2000	0.74%	62.30%	69.04%						
SLC	2000	1.42%	73.60%	80.12%	9.64%	8.43%	80.72%	81.93%	9.64%	9.64%
SMF	2000	0.65%	62.66%	65.90%	21.43%	38.46%	32.14%	19.23%	46.43%	42.31%
STL	2000	1.53%	69.14%	84.04%	4.55%	52.87%	79.55%	34.48%	15.91%	12.64%
Mean	2000.23	1.45%	64.77%	71.46%	19.79%	22.98%	55.02%	51.29%	25.19%	25.72%

LCC: low-cost carriers. In this and subsequent tables, data are missing for the empty cells.

However, covered airports put forth effort in a variety of other ways to increase competition. For example, both Charlotte Douglas International Airport (CLT) and San Antonio International Airport (SAT) implemented a marketing plan to attract additional low-fare carrier service. To make more efficient use of existing common-use facilities, ATL now enforces maximum turnaround times. Oakland International Airport (OAK) installed common use ticketing equipment (CUTE) at ticket counters and gates so that all airlines operating there will use identical facilities, providing maximum flexibility to airport administrators. CLT reduced landing fees for nonsignatory and new entrant carriers to the same level as signatory airlines (i.e., those obligated to repay debt issued by the airport). Nearly all covered airports implemented measures to record gate utilization, impose minimum-use standards, and notify airlines of gate availability to make more efficient use of existing gates. Many airports also amended majority-in-interest (MII) agreements to exempt capital projects necessary for competition from MII votes.[8]

III. Data

A. Sources

The majority of our data for this study are taken from the Data Bank 1B (DB1B) of the U.S. Department of

[8] MII agreements share the rights to decide on expansion of airport facilities between the airport authority and the airline controlling the majority of operations at the airport. In some cases, airlines even have veto power over airport expansions.

TABLE 2.—ENPLANEMENTS AND GATES FOR NONCOVERED AIRPORTS

	Enplanements			Gates					
	United States %	Top 2%		Common %		Legacy %		LCC %	
Airport	Mean	Mean	Maximum	2001	2008	2001	2008	2001	2008
BDL	0.49%	44.91%	49.15%						
BOS	1.87%	34.99%	38.93%	11.90%	10.31%	82.14%	68.04%	5.95%	21.65%
BUF	0.32%	37.59%	49.72%	6.25%	21.74%	78.13%	56.52%	15.63%	21.74%
CMH	0.49%	32.06%	37.66%	19.44%	22.22%	50.00%	58.33%	30.56%	19.44%
FLL	1.27%	35.05%	40.89%						
HNL	1.51%	45.34%	48.07%						
IND	0.57%	28.50%	32.86%	26.47%	30.00%	52.94%	57.50%	20.59%	12.50%
JFK	2.50%	41.42%	46.15%						
LAX	4.22%	34.44%	40.44%						
LGA	1.79%	41.34%	44.43%						
MCI	0.81%	42.96%	47.56%						
MCO	2.16%	36.66%	42.92%						
MSY	0.67%	44.68%	47.59%						
OKC	0.24%	41.82%	47.67%	0.00%	23.53%	68.75%	52.94%	31.25%	23.53%
OMA	0.28%	39.14%	41.72%	25.00%	35.00%	45.00%	45.00%	30.00%	20.00%
PDX	0.97%	37.15%	38.80%	19.57%	39.13%	32.61%	28.26%	47.83%	32.61%
RDU	0.63%	35.02%	41.34%	2.08%	19.05%	85.42%	66.67%	12.50%	14.29%
RSW	0.43%	39.12%	47.35%	23.53%	39.29%	58.82%	32.14%	17.65%	28.57%
SAN	1.15%	46.34%	47.80%	32.50%	22.50%	42.50%	43.75%	25.00%	33.75%
SEA	2.04%	45.12%	48.33%	21.62%	35.00%	40.54%	25.00%	37.84%	40.00%
SNA	0.62%	36.73%	39.64%						
TPA	1.20%	40.40%	42.66%	18.37%	28.81%	63.27%	44.07%	18.37%	27.12%
Mean	1.19%	39.13%	43.71%	17.23%	27.21%	58.34%	48.19%	24.43%	24.60%

Transportation's Origin and Destination Survey for the years 1993 through 2008. The DB1B data are a 10% random sample of all domestic itineraries. The unit of observation is the passenger level. The data contain information on the ticketing and operating carrier, details of any connections made by the passenger, and the fare paid for the itinerary used by the passenger. Following Evans and Kessides (1994), we consider round-trip tickets to be two equally priced one-way tickets and drop both interline and open-jaw tickets. Due to keypunch errors or redemption of frequent flier miles, there are some unusually large and small ticket prices in the DB1B data. For this reason, we drop any fares, measured in 2008 dollars, greater than $2,500 and less than $25.[9] In addition, we drop itineraries with more than six coupons (four connections) for round-trip itineraries and three coupons (two connections) for one-way itineraries. Following Borenstein and Rose (1994), we define a market as directional travel between a unique airport pair.

We supplement the DB1B data with information on the frequency and severity of delays from the DOT's Airline On-Time Performance database.

We also collected the enplanement data used by the FAA to determine coverage by AIR-21. There are significant differences between these data and the enplanement data that are publicly available through the DOT's T100 database. These differences arise because the T100 data do not include on-demand (e.g., charter flights) and in-transit (e.g., plane stops to refuel but does not deplane) passengers who are a significant source of enplanements at many airports. The differences are significant enough that the determination of coverage for a handful of airports would change depending on the source of enplanement data.

Our final source of data is a survey conducted jointly with the ACI-NA.[10] The survey, completed by 47% of all medium and large hubs (those enplaning more than 0.25% of all enplanements at primary airports in the United States), focused on gathering information on carrier-airport-specific leasing agreements for boarding gates. For each airport, we observe the total number of gates, number of gates leased by each carrier on an exclusive and preferential basis, and the number of gates reserved for common use by the airport authority. The details regarding the construction of the data set can be found in the online appendix.

B. *Descriptive Statistics*

Tables 1 and 2 summarize the FAA and survey data for medium and large hubs. Column 1 in tables 1 and 2 lists the covered and noncovered airports, respectively. The second column of table 1 lists the year in which each airport was first covered by the legislation. Due to the lag in data collection, coverage in any particular year is determined by enplanement data from two years earlier. For example, the set of airports first covered by the legislation in 2000 was determined using enplanement data from 1998. This is important for our purposes, since it would be very unlikely that an airline could perfectly foresee the details of the legislation two years in advance and manipulate enplanements to avoid coverage of a particular airport. Of covered airports, LAS was

[9] We also drop all itineraries for which the DOT questions the credibility of the reported fare, as indicated by the tktdollarcred variable.

[10] See Williams (2012) for more details on this survey.

the only airport not covered retroactively by the legislation. In section IVB, we test whether the lack of a significant number of airports first covered in later years is due to potential manipulation of enplanements by carriers.

The next three columns of tables 1 and 2 report the mean fraction of all U.S. enplanements performed at the airport, and the mean and maximum share of the top two carriers from 1998 to 2006 (determines coverage from 2000 to 2008). The maximum of the top two carriers' shares during this period serves as the predictor of coverage by the legislation since once an airport's concentration exceeds this threshold, it is required to file and execute a competition plan. Thus, for each airport in table 1 (2) this variable is greater (less) than .5. It is also important to note that coverage is not a proxy for the size of the airport. Examining the means at the bottom of tables 1 and 2 for the fraction of all U.S. enplanements, we see little difference in size between covered and noncovered airports. This is important, as it alleviates some concerns over the homogeneity of the treatment and control groups in our analysis.

The final columns of tables 1 and 2 report the fraction of gates reserved by the airport authority for common use, fraction leased on a preferential or exclusive basis by legacy carriers, and fraction leased on a preferential or exclusive basis by low-cost carriers (LCC).[11] Examining the respective means in 2001 and 2008 of these variables at the bottom of tables 1 and 2, we see little evidence that gates moved differentially at covered and noncovered airports. However, the large amount of missing data makes drawing any strong conclusions difficult. The lack of significant movement in the allocation of gates for most airports from 2001 to 2008 suggests that the FAA and DOT largely followed the recommendations put forth by GAO (2001). GAO (2001) cautioned that AIR-21 should not be used as a means to force the divestiture of assets (e.g., boarding gates) from dominant carriers at an airport for two reasons. First, the reallocation of assets among competing carriers might have little to no benefit if the gates were not allocated to a low-cost competitor (see Brueckner, Lee, & Singer, 2011, and Ciliberto & Tamer, 2009, for strong support for this statement). Second, service in smaller markets would likely be the first affected by divestiture of a dominant carrier's assets. This is intuitive: we expect a firm to eliminate or cut service in the least profitable markets, and the presence of any fixed costs associated with serving a market and significant economies of density in the industry (see Brueckner & Spiller (1994)) ensures a strong correlation between profitability and size.

The lack of a significant difference in the reallocation of gates among carriers at covered and noncovered airports foreshadows our finding that coverage by AIR-21 has little effect on the network of destinations offered out of an airport. It also suggests that if we are to find a significant effect from coverage by AIR-21 on other dimensions of service, it is due to more efficient use of existing assets (the focus of most competition plans) rather than to a redistribution of assets among carriers.

Table 3 summarizes the variables we construct from the DB1B data and other sources, before and after AIR-21, separately for the set of covered and noncovered airports. Column 3 gives the means of each variable, for covered and noncovered airports, over the entire sample period. To motivate our approach in section IV and emphasize the importance of controlling for trends in the data prior to coverage by AIR-21, we summarize the first difference for each variable. More precisely, for each variable, the difference before AIR-21 is calculated as the level in the first quarter of 1999 minus the level in the first quarter of 1993, while the difference after AIR-21 is calculated as the level in the first quarter of 2008 minus the level in the first quarter of 2002.[12]

Using the DB1B data, we classify a carrier's service in each market into one of two types: nonstop or connecting. For each type of service in a market, $Avg.Fare$ is calculated as the average fare across passengers choosing a type of service. $Avg.Mkt.Fare$, 20th$Pct.Mkt.Fare$, 50th$Pct.Mkt.Fare$, and 80th$Pct.Mkt.Fare$ are constructed similarly for different quantiles of the fare distribution in the market as a whole, aggregating across carriers and types of service. Table 3 shows a significant downward trend in fares in both covered and noncovered markets. However, prior to AIR-21, fares were falling less rapidly at covered airports, while after AIR-21, fares fell more rapidly at covered airports. These differential trends are strongest in the upper quantiles of the fare distribution. In section IV, we attempt to identify a causal relationship between coverage by AIR-21 and these differential trends in fares while controlling for a variety of time-varying covariates. $Nonstop$ is an indicator for whether a carrier's service is nonstop. $Distance.Traveled$ is the average number of miles traveled by passengers purchasing a type of service from a carrier in a particular market. For nonstop service, $Distance.Traveled$ is equal to the direct distance between the market end points. For connecting service, $Distance.Traveled$ is strictly greater than the direct distance.

At the airport-carrier level, we construct two variables from the DB1B. $Fraction.Routes$ is the proportion of all the destinations offered out of an originating airport for which a carrier offers either nonstop or connecting service. For example, suppose that 200 different markets are served out of Atlanta (ATL) by all carriers. Further, suppose Delta and American serve 180 and 100 markets out of ATL, respectively. In this example, $Fraction.Routes$ is equal to 0.9 and 0.5 for Delta and American, respectively. Thus, the numerator of $Fraction.Routes$ is the total number of markets served by a carrier out of an airport, $Number.Routes$, while the

[11] Low-cost carriers are B6, FL, F9, G4, J7, KP, KN, N7, NJ, NK, P9, QQ, SY, SX, TZ, U5, VX, W7, W9, WN, WV, XP, and ZA. The remaining carriers in our analysis are AA, AQ, AS, CO, DL, HA, NW, UA, US, and YX.

[12] Our results are nearly identical if 1993 to 2000 and 2001 to 2008 are used to construct the differences. Yet by not using 2000 and 2001, we reduce the possibility that our results are biased from the disruptions to the industry from the events of 9/11, anticipation by airports regarding the components of the law, or delays in its implementation.

TABLE 3.—MEANS FOR COVERED AND NONCOVERED MARKETS

	Covered	Mean, 1993–2008	Number of Observations	Pre-AIR21 1993–1999 Mean Diff.	Post-AIR21 2002–2008 Mean Diff.	Diff. in Diff.
Market-carrier-product						
Avg.Fare	Yes	234.926	7,669	−34.004	−21.695	−40.295
	No	274.505	1,353	−47.050	5.554	
Distance.Traveled	Yes	1.660	7,669	0.012	0.014	0.009
(Unit = 1000s of Miles)	No	1.658	1,353	0.020	0.013	
Fraction.Routes	Yes	0.380	7,669	−0.014	0.036	−0.004
	No	0.359	1,353	−0.010	0.044	
Pass	Yes	3,049.413	7,669	1,391.059	448.477	736.864
	No	2,688.051	1,353	1,152.424	−527.022	
Market						
Avg.Mkt.Fare	Yes	240.872	2,619	−55.993	−31.328	−29.281
	No	280.646	360	−63.127	−9.181	
20th Pct.Mkt.Fare	Yes	144.692	2,619	−47.002	0.704	−8.057
	No	168.967	360	−51.803	3.960	
50th Pct.Mkt.Fare	Yes	193.199	2,619	−66.638	−7.032	−18.589
	No	222.056	360	−72.163	6.032	
80th Pct.Mkt.Fare	Yes	312.435	2,619	−79.894	−55.398	−43.735
	No	370.831	360	−81.410	−13.179	
Avg.Dominant.Premium	Yes	17.604	1,360	17.516	−18.576	−18.789
	No	14.724	163	7.367	−9.936	
Lcc.Presence	Yes	0.723	2,619	0.337	0.183	0.192
	No	0.494	360	0.464	0.118	
Number.Firms	Yes	6.069	2,619	0.688	0.339	−0.169
	No	6.021	360	0.804	0.624	
P(delay>15)	Yes	6.127	105	−0.004	0.098	0.015
	No	6.924	1,372	0.018	0.106	
Pct.Nonstop	Yes	0.512	2,619	0.025	0.041	−0.059
	No	0.459	360	−0.012	0.062	
Airport-carrier						
Number.Routes	Yes	67.699	130	4.683	3.077	−3.401
	No	65.300	76	4.846	6.641	

denominator is the total number of routes served by any carrier out of the airport. This variable is intended to measure the relative attractiveness of carriers' networks.

Our measure of the hub premium, similar to the relative price measures used in Borenstein (1989), in a given market is calculated as the difference between the fares charged by the carriers with the largest share of enplanements at the origin and destination airports and the average of fares charged by all other carriers. For example, in the ATL (Atlanta Hartsfield) to CLE (Cleveland Hopkins) market, Delta and Continental are regarded as the dominant carriers (those with the largest share of enplanements), and *Avg.Dominant.Premium* is calculated as the difference between the average fare charged by Delta and Continental and the average fare charged by all other carriers. This variable is summarized in table 3 and suggests that coverage is associated with a moderate decline in the premium charged by dominant carriers. To measure the availability of nonstop service, an important dimension to service quality, we calculate *Pct.Nonstop* as the percentage of passengers traveling nonstop in a market. We also construct two measures of competition in a market, *Lcc.Penetration* and *Number.Firms*. *Lcc.Penetration*, summarized in table 3, is an indicator for whether a low-cost carrier is present in the market. As has been well documented, low-cost carrier penetration has been steadily increasing over the previous decade and typically results in intense price competition. In section IV, and as the descriptives suggest, we show that in markets where one or both end-point airports are covered by AIR-21, the low-cost penetration rate is significantly higher as a result of coverage. *Number.Firms* is the total number of firms serving the market and is a commonly used measure of competition in the industrial organization literature (see Berry, 1992, and Ciliberto & Tamer, 2009).

We calculate $P(Delay \geq 15)$, using the Airline On-Time Performance data, as the proportion of flights delayed fifteen minutes or more in a quarter on a particular flight segment. In addition to those variables we construct from the DOT sources, we also collected data on both population and per capita income for each MSA from the Bureau of Economic Analysis to serve as controls throughout our analysis.

IV. Empirical Analysis

Our final sample includes data from all airports classified as a medium or large hub by the FAA (enplaning at least 0.25% of total domestic enplanements), including highly concentrated hubs such as Minneapolis and Dallas.[13] A legitimate concern here is that these highly concentrated airports are significantly different from the control group (noncovered airports) in both observable and unobservable ways. For

[13] Smaller airports, below the 0.25% of total domestic enplanements, are not eligible to be covered by AIR-21. An analysis including these smaller airports gives nearly identical results. These estimates are available from the authors on request.

FIGURE 1.—COVERAGE CUTOFFS

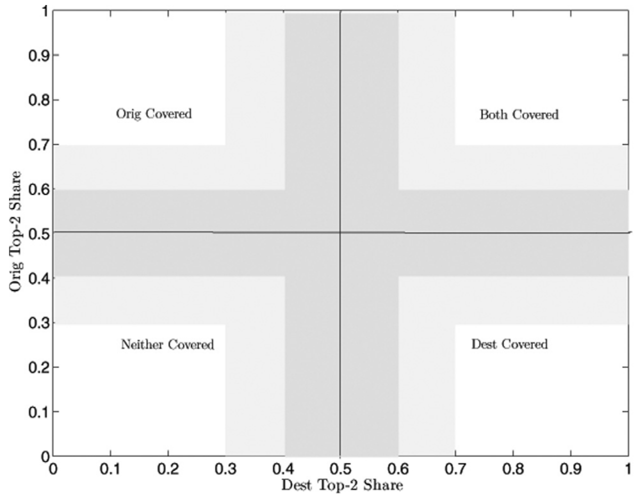

example, since airport presence is known to be an important factor in airline quality, cost, and price competition, it is troubling that we have no airports in the control group comparable to Dallas and Minneapolis in terms of presence measures. Similarly, unobserved airport features, such as geographic location, may affect the network economies of an airport, leading it to be highly concentrated and to have different competitive mechanics than less concentrated airports. The results from Berry and Jia (2010) also give an important example of the interaction of unobservable changes in consumer preferences (i.e., decreasing willingness to pay for quality), with observable airport presence differences.

To get around these problems, we exploit AIR-21's sharp discontinuity at the 50% two-carrier enplanement level. Essentially we assume that the distribution of market-level unobservables changes smoothly across the policy discontinuity. That is, the unobservable features of a randomly chosen market just below the cutoff have the same distribution as the unobservable features of a randomly chosen market just above the cutoff. Such an identification strategy also has the advantage of lessening concerns that other events (e.g., September 11) near the time of the policy intervention are influencing our results, as the event would have to have a differential effect on airports on either side of the treatment cutoff. A local identification approach also has the advantage of mitigating concerns regarding mean reversion, or the idea that concentrated airports tend to deconcentrate over time due to high prices and entry.

With this identification strategy in mind, we estimate the local average treatment effects (LATEs) of the law using two approaches. First, we proceed in the spirit of Black (1999) and estimate a series of difference-in-difference regressions using only those observations in progressively smaller windows around the concentration cutoffs determining coverage. Figure 1 demonstrates this approach. We begin by using the complete sample and then examine the subset of markets within narrower windows around the coverage cutoffs. Using this approach, we identify market outcomes affected by coverage in a statistically and economically significant manner. This approach also allows us to use covariates to control for observable differences in airports and markets. This is potentially useful because while we have a large number of markets, these markets are drawn from a relatively small number of airports, which may create a small sample problem even if our identifying assumption is correct. For example, New York (JFK) is always included as a control airport and serves destinations that are larger, richer, and more distant on average than those in the treatment group, and due to the large number of markets originating or terminating at the airport, it represents a nontrivial fraction of the sample.

In the top half of figure 2, we plot the cdf of per capita income and population of destinations served out of covered and noncovered airports. Airports covered by AIR-21 tend to serve destinations with larger populations and higher per capita incomes. In the bottom half of figure 2, we plot the cdf of the same observables for destinations served out of airports with two-firm concentrations within 0.1 of the treatment cutoff. In this window, the difference in the distribution of population narrows greatly, while the difference in the right-hand tails of the distribution of per capita income also narrows. Figure 3 shows the two-firm concentration ratios at each airport in 2000 and 2008, respectively.

In the second step, we employ a true RD approach and allow the window width to collapse to 0. We find that our main conclusions from the first step are robust, which one would expect given the similarity of the distribution of observables in figure 2. In addition, our particular RD implementation allows us to examine variation in the effect of coverage along the cutoffs. We do this by defining treatment effects as functions of an origin's or destination's concentration level, which we describe more fully below, whereas in the first step, we identify only an average effect of coverage, ignoring potential heterogeneity in treatment effects across markets. In the case of low-cost entry, we find significant heterogeneity along the treatment cutoff. Intuitively, following implementation of the competition plans and the resulting improvement in access to airport facilities, we find that low-cost carriers enter more concentrated markets with greater frequency.

A. *Boundary Dummy Approach*

Fares. Following Black (1999), we begin under the assumption that coverage is exogenous and homogeneous in its effect on fares by estimating the following regression:

$$\begin{aligned}\Delta_t \log(Avg.Fare_{ijmt}) &= \Delta_t x_{ijmt}\beta + \Delta_t z_{mt}\gamma + \psi Nonstop_{ijmt} \\ &\quad + \tau_1 1[1 \text{ cover}_m] + \tau_2 1[2 \text{ cover}_m] + \Delta_t \epsilon_{ijmt}, \qquad (1)\end{aligned}$$

using the complete sample. The dependent variable is the long second difference, the change from 2002 to 2008 minus the change from 1993 to 1999, of the logarithm of average fares paid by passengers who purchased product j (nonstop

FIGURE 2.—DISTRIBUTION OF INCOME AND POPULATION OF DESTINATIONS SERVED, COVERED VERSUS NONCOVERED AIRPORTS

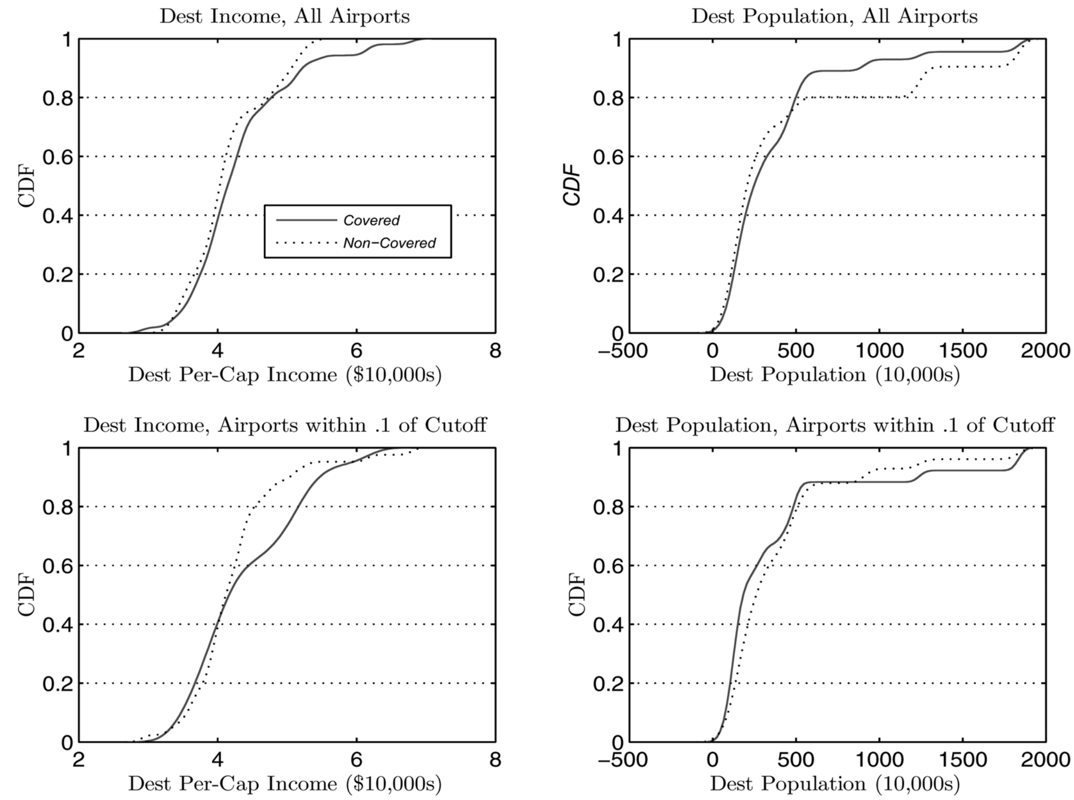

FIGURE 3.—CHANGE IN TOP TWO CONCENTRATIONS

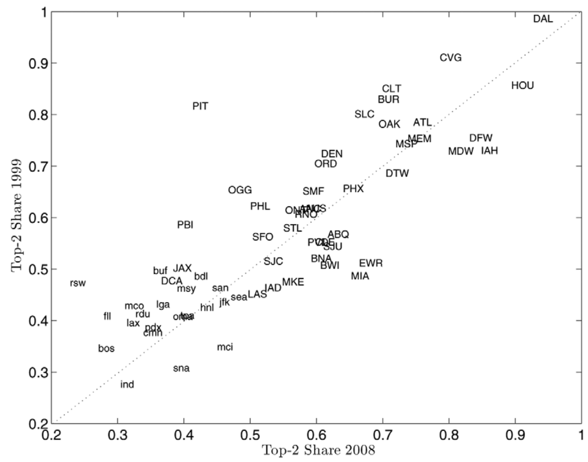

or connecting service) from carrier i in market m.[14] The vectors $\Delta_t x_{ijmt}$ and $\Delta_t z_{mt}$ include the second differences of *Fraction.Routes*, *Distance.Traveled*, and the population and per capita income at the market end point airports.[15] In addition, we include an indicator for nonstop service to capture the possibility that fares for nonstop service changed differentially relative to connecting service.

To capture the impact of coverage by AIR-21 on the time path of fares, we include indicators for whether one or both of a market's end points were covered, $1[1\ \text{cover}_m]$ and $1[2\ \text{cover}_m]$, respectively.[16] Under the assumption that coverage is exogenous and homogeneous in its effect on fares, τ_1 and τ_2 measure the causal effect on the dependent variable in a market with one and two end points covered. To relax these assumptions and ensure a causal interpretation of τ_1 and τ_2, we estimate the same regression on the subsamples of markets in progressively smaller windows around the coverage cutoffs. For such an approach to give consistent estimates, a significant portion of the data must be located within these windows. Figure 4a gives the relative mass of the number of observations for each combination of the predictors of treatment. Figure 4b provides a contour map of the density and demonstrates that the majority of the data is in fact immediately around the coverage cutoffs. This is of particular importance as we shrink the window further in the RD analysis.

The estimates of equation (1) are presented in columns 1, 3, and 5 of table 4. For the regression results presented in tables 4 to 9, we calculate robust standard errors using the

[14] The DB1B first became publicly available in 1993. The results are insensitive to varying the first or last year of the panel used in the analysis.

[15] See Berry (1990), Berry, Carnall, and Spiller (2006), and Berry and Jia (2010) for a discussion of the impact of the size of a carrier's network on demand for that carrier's services.

[16] In earlier versions of the paper, we included specifications that allowed the $1[1\ \text{cover}_m]$ treatment effect to vary according to whether the origin or destination is treated. The effects were highly symmetric, so we opted for the more parsimonious specification.

FIGURE 4.—MASS OVER SUPPORT OF PREDICTORS

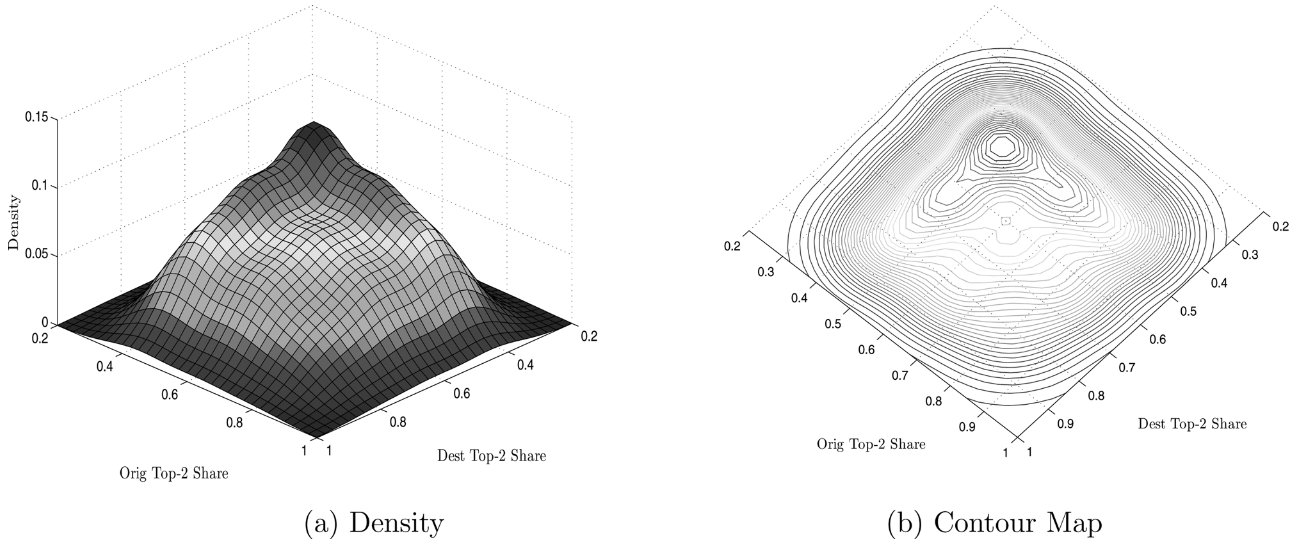

TABLE 4.—AVERAGE FARE REGRESSIONS

	All Markets		0.2 of Cutoff		0.1 of Cutoff	
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Avg. Fare)	$N = 8{,}996$		$N = 8{,}453$		$N = 5{,}813$	
1[1 cover]	−0.147***	−0.106***	−0.145***	−0.105***	−0.144***	−0.109***
	(0.030)	(0.023)	(0.030)	(0.023)	(0.030)	(0.024)
1[2 cover]	−0.210***	−0.133***	−0.202***	−0.129***	−0.226***	−0.143***
	(0.047)	(0.033)	(0.046)	(0.032)	(0.050)	(0.034)
Nonstop	−0.191***	−0.189***	−0.195***	−0.193***	−0.184***	−0.182***
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)
Fraction.Routes	0.275***	0.314***	0.277***	0.304***	0.291***	0.310***
	(0.104)	(0.079)	(0.100)	(0.078)	(0.097)	(0.081)
Distance.Traveled	0.161***	0.196***	0.152***	0.184***	0.125**	0.154***
	(0.047)	(0.041)	(0.049)	(0.044)	(0.055)	(0.046)
Lcc.Presence		−0.166***		−0.164***		−0.151***
		(0.020)		(0.020)		(0.022)
Number.Firms		−0.013***		−0.013**		−0.012**
		(0.005)		(0.005)		(0.006)
Origin.HHI		0.142		0.161		0.051
		(0.183)		(0.183)		(0.173)
Dest.HHI		0.046		0.066		−0.017
		(0.167)		(0.164)		(0.157)
Market.HHI		0.067**		0.074**		0.121***
		(0.034)		(0.034)		(0.040)
Has.Other.Products		0.007		0.009		0.005
		(0.007)		(0.008)		(0.008)
R^2	0.096	0.178	0.092	0.174	0.098	0.175
Borenstein-Rose (1994) controls	No	Yes	No	Yes	No	Yes

Additional controls include Population Origin, Population Dest, Per-Cap Income Origin, and Per-Cap Income Dest. Statistical significance: ***1%, **5%, *10%.

approach of Cameron, Gelbach, and Miller (2007) to cluster on both the origin and destination, allowing arbitrary forms of correlation between observations with either the origin, destination, or both in common. Our estimates of τ_1 and τ_2 are negative and statistically and economically significant in each window around the coverage cutoff. From column 5, where we can reasonably interpret our coefficients in a causal fashion, the results indicate that coverage of a single end point by AIR-21 results in a 13.4% reduction in average fares,[17] while coverage of both end points results in approximately a 20.2% change in average fares. This result is robust across different window widths, suggesting we might safely extrapolate our local effects to more highly concentrated airports. The remaining results in column 5 are straightforward to interpret and are consistent across all subsamples.

We interpret the results in columns 1, 3, and 5 as estimates of the total effect, both direct and indirect, of AIR-21 on average fares. Since we typically think of the effect of barriers to entry on fares as being indirect, that is, barriers to entry affect entry and exit, and market structure, which in turn affect prices, it is important to the credibility of our

[17] Following Halvorsen and Palmquist (1980), this is calculated as $-13.4 = 100 \times (\exp(-0.144) - 1)$.

identification as well as our evaluation of the policy to try to understand the direct channels through which fares are affected. Columns 2, 4, and 6 of table 4 show specifications designed to partially get at these channels. Specifically, we add a very similar set of time-varying regressors to those employed by Borenstein and Rose (1994) to control for any changes in the competitive environment in a market. The set of controls includes the number of competitors in the market, an indicator for whether a low-cost carrier serves the market, airport-level Herfindahl indices for both end points, market-level Herfindahl indices, and an indicator for whether the carrier offers both nonstop and connecting service.[18] Each of these variables enters the regression as a second difference as described above. One can then interpret changes in the estimates of τ_1 and τ_2, when the controls are included, as evidence that variation in the competitive environment explains some portion of the estimated effect of coverage. We find that these controls explain between 30% and 50%, depending on the window width, of the effect from coverage we estimated in columns 1, 3, and 5. *Lcc.Penetration* has a statistically significant and economically meaningful effect on fares. This suggests the explanation for the large decline in fares caused by AIR-21 has come, to a large extent, through exactly the channels envisioned by the law. While, under our identifying assumptions, our treatment effect estimates are consistent in the presence of potentially endogenous market structure variables, we do not argue that the coefficients on these variables are consistent estimates of the effect of, for example, low-cost entry. However, we note that the substantial 14% declines associated with the *Lcc.Penetration* variable are consistent with the range of estimates presented by Brueckner et al. (2011).[19]

Estimates of equation (1) provide insight into how fares for a particular type of service in market changed as a result of AIR-21. It is also of interest to know how the distribution of fares in a market as a whole changed as a result of AIR-21. Specifically, performing the analysis on the distribution of fares in the market as a whole serves as an informal test of whether sample selection, on carriers serving a market over the entire period, is driving our results. If the estimated effects for coverage are similar to those of equation (1), we can be confident that selection is not driving our results.

For this purpose, we estimate market-level regressions of the form

$$\begin{aligned}\Delta_t \log(Qtile.Mkt.Fare_{mt}) \\ = \Delta_t z_{mt}\gamma + \tau_1 1[1\ \text{cover}_m] + \tau_2 1[2\ \text{cover}_m] + \Delta_t \epsilon_{mt},\end{aligned}$$

where the dependent variable, $Qtile.Mkt.Fare_{mt}$, is a particular quantile (e.g., 20th*Pct.Mkt.Fare*) of the distribution of fares in a market. Table 5 presents these estimates. For conciseness and due to the similarity of the estimates to those in table 4, we present only coefficients of particular interest. Columns 2, 4, and 6 (1, 3, and 5) present the estimates with(out) the Borenstein and Rose (1994) controls; the number of competitors in the market, an indicator for whether a low-cost carrier serves the market, airport-level enplanement Herfindahl indices for both end points, and market-level enplanement Herfindahl indices.

The results for $Avg.Mkt.Fare_{mt}$, in the top panel of table 5, are very similar to those reported in table 4, suggesting that selection is not a concern. Additionally, consistent with the descriptive evidence in table 3, we find that the estimated decline in fares resulting from coverage by AIR-21 is increasing in the fare quantile. The coefficients in column 5, which we can reasonably interpret in a causal fashion, show little decline at the bottom of the distribution as a result of coverage, while there were large and statistically significant declines in both 50th*Pct.Mkt.Fare* and 80th*Pct.Mkt.Fare*. 50th*Pct.Mkt.Fare* declines 8.6% (19.0%) when one (both) end point(s) is (are) covered, while 80th*Pct.Mkt.Fare* declines 10.1% (21.6%). Inclusion of the Borenstein and Rose (1994) controls, low-cost carrier entry in particular, again explains away up to 50% of the estimated effect from coverage. Consistent with low-cost carriers successfully targeting business travelers, we find the effect of low-cost entry to be larger in the upper quantiles of the fare distribution.

The last measure of the impact of AIR-21 on fares that we examine is the effect on the dominant carrier premium, $Avg.Dominant.Premium_{mt}$. We measure the hub premium as the difference in the logarithm of the average fare charged in a market by the carriers with the largest presence at the market's end points with that of its competitors. These premiums range from roughly 15% to 40% in 1999 and, on average, are sharply increasing in the concentration of an airport. Table 6 reports the results of these regressions. The magnitudes of the estimates are similar across different window widths but not always statistically significant. If we focus on column 5, the narrowest window, we find that these premiums have fallen slightly faster in markets with one or both end points covered, although the coefficient on 1[2 cover] is not statistically significant. These results suggest, at most, that AIR-21 had a modest impact on reducing premiums charged by dominant carriers, despite the success of AIR-21 in reducing fares overall.

Quality. In addition to fares, many other characteristics of service may change as the result of coverage by AIR-21. GAO (2001) suggests that granting authority to regulators to force dominant carriers at certain airports to divest critical assets (e.g., boarding gates) introduces uncertainty and can lead to disinvestment in an airport. In particular, GAO (2001) suggests that service to smaller airports would be the first to be affected, possibly losing service altogether. If the fare reductions are accompanied by diminished service quality

[18] We also estimated a specification including an interaction of the nonstop indicator and the indicator for whether the carrier offers both nonstop and connecting service. This interaction was never significant.

[19] Results from estimating slight variations of equation (1) using the entire panel of data (every year from 1993 to 2008) and year-on-year first differences are available at http://jonwms.myweb.uga.edu/research.htm. The results are nearly identical to those presented in table 4.

TABLE 5.—FARE DISTRIBUTION REGRESSIONS

	All Markets		0.2 of Cutoff		0.1 of Cutoff	
	(1)	(2)	(3)	(4)	(5)	(6)
	$N = 2{,}978$		$N = 2{,}746$		$N = 1{,}825$	
Log(Avg.Mkt.Fare)						
1[1 cover]	−0.116***	−0.090***	−0.117***	−0.089***	−0.115***	−0.092***
	(0.034)	(0.026)	(0.034)	(0.026)	(0.032)	(0.027)
1[2 cover]	−0.213***	−0.139***	−0.198***	−0.128***	−0.238***	−0.152***
	(0.054)	(0.039)	(0.054)	(0.040)	(0.064)	(0.050)
Lcc.Presence		−0.257***		−0.253***		−0.227***
		(0.026)		(0.027)		(0.027)
Number.Firms		−0.009		−0.009		−0.003
		(0.007)		(0.007)		(0.008)
R^2	0.050	0.287	0.047	0.277	0.071	0.264
Log(20thPct.Mkt.Fare)						
1[1 cover]	−0.063*	−0.052*	−0.063*	−0.051*	−0.068**	−0.058**
	(0.034)	(0.029)	(0.034)	(0.029)	(0.033)	(0.028)
1[2 cover]	−0.088*	−0.058	−0.089*	−0.061	−0.078	−0.032
	(0.049)	(0.041)	(0.051)	(0.043)	(0.065)	(0.057)
Lcc.Presence		−0.120***		−0.118***		−0.108***
		(0.021)		(0.021)		(0.023)
Number.Firms		−0.022***		−0.022***		−0.019***
		(0.006)		(0.006)		(0.007)
R^2	0.008	0.109	0.009	0.110	0.014	0.096
Log(50th Pct.Mkt.Fare)						
1[1 cover]	−0.092***	−0.072***	−0.091***	−0.073***	−0.090***	−0.076***
	(0.032)	(0.025)	(0.032)	(0.025)	(0.033)	(0.028)
1[2 cover]	−0.199***	−0.145***	−0.186***	−0.139***	−0.211***	−0.143**
	(0.049)	(0.037)	(0.051)	(0.040)	(0.070)	(0.060)
Lcc.Presence		−0.208***		−0.197***		−0.166***
		(0.028)		(0.029)		(0.030)
Number.Firms		−0.030***		−0.031***		−0.023**
		(0.008)		(0.008)		(0.009)
R^2	0.031	0.184	0.028	0.177	0.037	0.151
Log(80thPct.Mkt.Fare)						
1[1 cover]	−0.113**	−0.077**	−0.114***	−0.077**	−0.107**	−0.077**
	(0.044)	(0.035)	(0.044)	(0.034)	(0.042)	(0.034)
1[2 cover]	−0.242***	−0.144***	−0.213***	−0.118**	−0.243***	−0.124*
	(0.071)	(0.053)	(0.070)	(0.051)	(0.083)	(0.064)
Lcc.Presence		−0.340***		−0.336***		−0.302***
		(0.030)		(0.031)		(0.033)
Number.Firms		−0.011		−0.011		−0.004
		(0.007)		(0.007)		(0.008)
R^2	0.046	0.251	0.041	0.242	0.062	0.242
Borenstein-Rose (1994) controls	No	Yes	No	Yes	No	Yes

The coefficients of the remaining Borenstein-Rose (1994) controls, Origin HHI, Dest HHI, and Market HHI, are not displayed for space considerations and are available from the authors on request. Additional controls include Population Origin, Population Dest, Per-Cap Income Origin, Per-Cap Income Dest. Statistical significance: ***1%, **5%, *10%.

TABLE 6.—DOMINANT CARRIER PREMIUM REGRESSIONS

	All Markets		0.2 of Cutoff		0.1 of Cutoff	
	(1)	(2)	(3)	(4)	(5)	(6)
	$N = 1{,}487$		$N = 1{,}343$		$N = 899$	
Avg.Dominant.Premium						
1[1 cover]	−0.070*	−0.064*	−0.068*	−0.062	−0.076**	−0.072*
	(0.037)	(0.038)	(0.036)	(0.038)	(0.037)	(0.039)
1[2 cover]	−0.063	−0.049	−0.070	−0.054	−0.038	−0.026
	(0.048)	(0.049)	(0.051)	(0.052)	(0.064)	(0.061)
Lcc.Presence		−0.029		−0.036		−0.021
		(0.026)		(0.026)		(0.023)
Number.Firms		0.002		0.005		0.005
		(0.008)		(0.009)		(0.009)
R^2	0.004	0.014	0.005	0.017	0.011	0.026
Borenstein-Rose (1994) controls	No	Yes	No	Yes	No	Yes

The coefficients of the remaining Borenstein-Rose (1994) controls, Origin HHI, Dest HHI, and Market HHI, are not displayed for space considerations and are available from the authors on request. Additional controls include Population Origin, Population Dest, Per-Cap Income Origin, Per-Cap Income Dest. Statistical significance: ***1%, **5%, *10%.

then the welfare consequences of coverage are ambiguous. We focus on three critical dimensions of service quality: the availability of nonstop service (percentage of passengers flying nonstop in a market), the on-time performance of carriers (fraction of flights delayed fifteen minutes of more), and the number of markets served by a carrier out of an airport (number of destinations served on a connecting or nonstop basis by a carrier out of an airport).

To estimate the impact of coverage on the availability of nonstop service, we estimate the following regression:

$$\begin{aligned}\Delta_t Pct.Nonstop_{mt} \\ &= \Delta_t z_{mt}\gamma + \tau_1 1[1 \text{ cover}_m] + \tau_2 1[2 \text{ cover}_m] + \Delta_t \epsilon_{mt},\end{aligned}$$

where $\Delta_t Pct.Nonstop_{mt}$ denotes the second difference, constructed identically to the dependent variable in equation (1), in the fraction of passengers flying nonstop in market m. To examine the impact of coverage on the frequency of late flights on any nonstop flight segment s, we estimate the following regression:

$$\begin{aligned}\Delta_t P(Delay \geq 15)_{st} \\ &= \Delta_t z_{st}\gamma + \tau_1 1[1 \text{ cover}_m] + \tau_2 1[2 \text{ cover}_m] + \Delta_t \epsilon_{st},\end{aligned}$$

where $P(Delay \geq 15)$ is the fraction of flights delayed fifteen minutes or more. Finally, to capture any potential divestiture by carriers in an airport resulting from coverage by AIR-21, we estimate the following regression:

$$\begin{aligned}\Delta_t \log(Number.Routes_{iat}) \\ &= \Delta_t z_{at}\gamma + \tau 1[1 \text{ cover}_a] + \Delta_t \epsilon_{iat},\end{aligned}$$

where the unit of observation is at the carrier (i)-airport (a) level.

The results of these regressions are presented in table 7. The GAO (2001) was most concerned that AIR-21 might lead carriers to completely cease operations or reduce the number of destinations served from covered airports. We do not find robust evidence that the range of destinations offered by carriers was affected negatively by coverage; however, the large amount of noise in these regressions and the variation in the point estimates make it difficult to draw firm conclusions. With regard to delays, consistent with the findings of Forbes (2008), we find weak evidence of an increase in the proportion of late flights; however, this result is not robust across all windows.[20] This is an intuitive result. Mayer and Sinai (2003) find that carriers controlling the majority of the operations at an airport have an incentive to internalize congestion-related delays. At covered airports, a number of gates were seized from dominant carriers and subsequently operated by the airport authority, a third party with a lesser incentive to minimize delays of any particular carrier or flight. Besides an increase in delays, we find that the proportion of passengers flying nonstop declines slightly as a result of coverage. While this is difficult to test, a likely explanation is a change in the composition of passengers in covered markets, such that those consumers brought into the market as a result of the large declines in fares are more price sensitive and prefer cheaper connecting service.

Table 7.—Quality Regressions

	All Markets (1)	0.2 of Cutoff (2)	0.1 of Cutoff (3)
Pct.Nonstop	$N = 2{,}978$	$N = 2{,}746$	$N = 1{,}825$
1[1 cover]	−0.052**	−0.052**	−0.045*
	(0.026)	(0.026)	(0.024)
1[2 cover]	−0.063**	−0.064**	−0.038
	(0.030)	(0.030)	(0.032)
R^2	0.007	0.007	0.009
P(Delay>15)	$N = 1{,}420$	$N = 1{,}263$	$N = 814$
1[1 cover]	0.038*	0.039*	0.035*
	(0.020)	(0.020)	(0.020)
1[2 cover]	0.067**	0.064**	0.045
	(0.027)	(0.027)	(0.031)
R^2	0.074	0.083	0.097
Log(Number.Routes)	$N = 419$	$N = 312$	$N = 181$
1[cover]	−0.069	−0.092*	−0.010
	(0.054)	(0.055)	(0.082)
R^2	0.008	0.017	0.011

Additional controls include Population Origin, Population Dest, Per-Cap Income Origin, Per-Cap Income Dest. Statistical significance: ***1%, **5%, *10%.

If it were not for such a significant decline in fares, conclusions regarding consumer welfare would be—and to a significant degree still are—less clear, as it would require an unreasonably high opportunity cost of time (associated with increased delays and connections) to completely offset a 19% reduction in the median fare. However, this finding speaks to the idea that AIR-21 may have been suboptimally focused on the reallocation of scarce facilities rather than expansion of the facilities.

Quantity. For the decrease in fares in covered markets to be welfare improving for consumers, we should not observe statistically significant declines in passenger traffic in these markets. We have shown there are only minor declines in measurable aspects of service quality. Observing declines in passenger volumes would suggest that other, unobserved, aspects of service (e.g., quality) declined along with fares in covered markets. To measure the effect of coverage on passenger volumes, we run the following regression:

$$\begin{aligned}\Delta_t \log(Pass_{ijmt}) = \Delta_t x_{ijmt}\beta + \Delta_t z_{mt}\gamma + \psi Nonstop_{ijmt} \\ + \tau_1 1[1 \text{ cover}_m] + \tau_2 1[2 \text{ cover}_m] \\ + \Delta_t \epsilon_{ijmt},\end{aligned}$$

where $Pass_{ijmt}$ is the number of passengers traveling, either nonstop or connecting, j, with carrier i in market m, as reported in the DB1B data. The vectors $\Delta_t x_{ijmt}$ and $\Delta_t z_{mt}$ include the same controls as in equation (1).

[20] In unreported results, we included a variety of measures to attempt to measure congestion and concentration to try to explain the observed increase in the frequency of late flights. Interestingly, none of these measures explain away any portion of the measured effect of coverage. This leaves very few possible explanations other than a change in control over the management of the scarce airport facilities.

TABLE 8.—QUANTITY REGRESSIONS

	All Markets		0.2 of Cutoff		0.1 of Cutoff	
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Pass)	$N = 8{,}724$		$N = 8{,}194$		$N = 5{,}650$	
1[1 cover]	0.199***	0.076	0.200***	0.074	0.185***	0.059
	(0.044)	(0.047)	(0.044)	(0.047)	(0.043)	(0.050)
1[2 cover]	0.267***	0.027	0.264***	0.022	0.352***	0.138
	(0.081)	(0.076)	(0.085)	(0.078)	(0.115)	(0.103)
Lcc.Presence		0.268***		0.271***		0.254***
		(0.030)		(0.032)		(0.030)
Number.Firms		0.032***		0.031***		0.029***
		(0.011)		(0.011)		(0.010)
R^2	0.059	0.267	0.060	0.270	0.061	0.280
Borenstein-Rose (1994) controls	No	Yes	No	Yes	No	Yes

The coefficients of the remaining Borenstein-Rose (1994) controls, Origin HHI, Dest HHI, and Market HHI, are not displayed for space considerations and are available from the authors on request. Additional controls include Population Origin, Population Dest, Per-Cap Income Origin, Per-Cap Income Dest. Statistical significance: ***1%, **5%, *10%.

The results of these regressions are reported in table 8. Columns 2, 4, and 6 (1, 3, and 5) report the results with(out) the Borenstein and Rose (1994) controls for different window widths. Column 5 indicates that coverage, on average, had a positive and statistically significant effect on passenger volumes. Passenger volumes rose approximately 20.3% (42.2%) in markets with one (both) end point(s) covered, consistent with the large coverage effect for fares reported in table 4. The large increase in passenger volumes suggests that the increase in delays did little to diminish the benefit to consumers resulting from the large decline in fares. Column 6 shows that inclusion of the Borenstein and Rose (1994) controls explains away much of the estimated effect on passenger volumes from coverage. Low-cost entry in covered markets drives nearly all of the increase in passenger volumes.

TABLE 9.—COMPETITION REGRESSIONS

	All Markets (1)	0.2 of Cutoff (2)	0.1 of Cutoff (3)
Lcc.Presence	$N = 2{,}978$	$N = 2{,}746$	$N = 1{,}825$
1[1 cover]	0.114***	0.117***	0.108**
	(0.043)	(0.044)	(0.043)
1[2 cover]	0.305***	0.298***	0.404***
	(0.084)	(0.084)	(0.106)
R^2	0.008	0.008	0.020
Log(Number.Firms)	$N = 2{,}978$	$N = 2{,}746$	$N = 1{,}825$
1[1 cover]	−0.142	−0.137	−0.105
	(0.214)	(0.215)	(0.223)
1[2 cover]	−0.351	−0.377	−0.053
	(0.325)	(0.313)	(0.330)
R^2	0.006	0.007	0.007

Additional controls include Population Origin, Population Dest, Per-Cap Income Origin, Per-Cap Income Dest. Statistical significance: ***1%, **5%, *10%.

Competition. The results in tables 4, 5, and 8 suggest that increased competition, particularly by low-cost carriers, explains a significant portion of the variation in fares and increases in passenger volumes in covered markets. However, it is not clear whether this increase in competition is driven by coverage. To test whether the steps taken by covered airports had a significant impact on the number and identity of firms, we estimate two regressions:

$$\Delta_t Number.Firms_{mt} = \Delta_t z_{mt}\gamma + \tau_1 1[1\text{ cover}_m] + \tau_2 1[2\text{ cover}_m] + \Delta\epsilon_{mt}$$

and

$$\Delta_t Lcc.Penetration_{mt} = \Delta_t z_{mt}\gamma + \tau_1 1[1\text{ cover}_m] + \tau_2 1[2\text{ cover}_m] + \Delta\epsilon_{mt},$$

where the dependent variables in these regressions are the number of firms serving the market and an indicator for whether a low-cost carrier is present, respectively.

The estimates of the coefficients on the coverage indicators are presented in table 9. We find that for markets with one (both) end point(s) covered, there is a 0.11 (0.40) increase in the probability of a low-cost carrier serving the market. While there are obvious caveats in interpreting the coefficients of a second differenced linear probability model, at a minimum, this corroborates our finding that variation in the low-cost indicator played a major role in explaining between 30% and 50% of the reduction in fares as a result of coverage. We find AIR-21 has no significant impact on the average number of firms serving a market.

B. Regression Discontinuity Design

There are many strengths associated with Black's (1999) approach. The results, however, rely on a number of assumptions, including homogeneity of the coverage effect and exogeneity of coverage, to estimate the effects of coverage by AIR-21. These assumptions can be troublesome because more concentrated airports (those with two carriers enplaning more than 50% of the passengers) are treated, while less concentrated airports are not. Therefore, any covariation between fares and concentration after the first quarter of 2001 (the time of the treatment) would be empirically indistinguishable from a treatment effect due to AIR-21. While these assumptions are difficult to test formally, it is possible to measure a local-average treatment effect (LATE) around the treatment cutoff in the absence of these assumptions using an RD approach. Examining treatment and control groups "very near" either side of the treatment cutoff allows us to disentangle those movements in fares that result from coverage from those that are simply due to correlation between fares and concentration. We discuss our approach below.[21]

Defining LATEs. Estimation of the LATEs here is complicated by the two-dimensional predictor vector. Instead of a point, our LATE estimates are now functions of the predictors of treatment, P_m^{orig} and P_m^{dest}. Let $Y_{ijmt}(o,d)$, $o,d \in \{0,1\}$, denote the outcome variable when the origin treatment status is o and the destination treatment status is d. For each observation, we get to observe one of the four possible values of the variable. When only one end point is treated, we define the LATEs as

$$\tau_1^{orig}(P_m^{dest}) = E[Y_{imt}(1,0) - Y_{imt}(0,0) | P_m^{orig} = .5 P_m^{dest} < .5],$$
$$\tau_1^{dest}(P_m^{orig}) = E[Y_{imt}(0,1) - Y_{imt}(0,0) | P_m^{orig} < .5 P_m^{dest} = .5],$$

and when both end points are treated,

$$\tau_2^{orig}(P_m^{dest}) = E[Y_{imt}(1,1) - Y_{imt}(0,1) | P_m^{orig} = .5 P_m^{dest} > .5],$$
$$\tau_2^{dest}(P_m^{orig}) = E[Y_{imt}(1,1) - Y_{imt}(1,0) | P_m^{orig} > .5 P_m^{dest} = .5].$$

Our definition of treatment effects is motivated by several considerations. First are identification considerations. Our data are lumpy in the sense that the predictors of coverage do not vary within an airport, so for a sufficiently small window around a given concentration level, all the markets in that window will be drawn from a single airport. For example, consider Dallas–Fort Worth (DFW), which has a predictor value of around 0.8, well away from the coverage cutoff. The estimate of $\tau_{dest}^2(0.8)$ compares the path of fares over the period since the passage of AIR-21 in markets originating at DFW and terminating at airports just below the coverage cutoff to markets originating at DFW and terminating at airports just above the coverage cutoff. This approach allows us to control, to some extent, for fixed unobserved factors associated with given airports that are potentially distant from the coverage cutoffs. Second, in contrast to the window regressions, allowing the treatment effect to vary along the treatment cutoff in addition to the local linear regression implementation, discussed below, we estimate the effect of coverage more flexibly. Figure 4 shows that the majority of the observations are near the treatment cutoff, making such a flexible approach feasible. Moreover, Berry and Jia (2010) suggest there is direct evidence that the treatment effects may differ in airport concentration. Of course, the interpretation of our estimates as a flexible interactive effect is invalid if there is selection inherent in conditioning on the away-from-the-boundary-airport concentration level, which is likely given that a single airport will dominate any small bin. However, even in the presence of such selection, we can still interpret the estimates as an estimate of LATE heterogeneity where the heterogeneity corresponds to interaction with whatever is driving selection.

Our major task in estimation is to adapt the basic RD framework to account for a two-dimensional predictor vector. This requires flexibly estimating a two-dimensional surface that relates Y_{imt} to $\{P_m^{orig}, P_m^{dest}\}$. Local linear estimators are particularly attractive for this type of problem (see Imbens & Lemieux, 2008). At boundary points of the support for the predictor vector, local linear estimators do not suffer from the inherent bias of kernel estimators and achieve faster rates of convergence. In addition, local linear estimators are easily extended to multiple dimensions. Fan and Gijbels (1996) provide a detailed discussion of the advantages of local-polynomial modeling.

To demonstrate our approach, suppose we are estimating $\tau_1^{orig}(P_m^{dest})$. This requires us to estimate the conditional expectation, $E[Y_{imt}(1,0) - Y_{imt}(0,0) | P_m^{orig} = .5, P_m^{dest} < .5]$, for each $P_m^{dest} < .5$. For a particular value of P_m^{dest}, $\overline{P}^{dest}$, the estimator is defined as

$$\tau_1^{orig}(\overline{P}^{dest}) = \widehat{\alpha}^{c+} - \widehat{\alpha}^{c-},$$

where

$$\min_{\left\{\alpha^{c-},\beta^{c-}_{orig},\beta^{c-}_{dest}\right\}} \sum_{\substack{m:\\ \{P_m^{orig} < .5, P_m^{dest} < .5\}}} [Y_{imt}(0,0) - \alpha_0^{c-} - \beta_{orig}^{c-}(P_m^{orig} - .5) - \beta_{dest}^{c-}(P_m^{dest} - \overline{P}^{dest})]^2 w_m^- \tag{2}$$

and

$$\min_{\left\{\alpha^{c+},\beta^{c+}_{orig},\beta^{c+}_{dest}\right\}} \sum_{\substack{m:\\ \{P_m^{orig} \geq .5, P_m^{dest} < .5\}}} [Y_{imt}(1,0) - \alpha_0^{c+} - \beta_{orig}^{c+}(P_m^{orig} - .5) - \beta_{dest}^{c+}(P_m^{dest} - \overline{P}^{dest})]^2 w_m^+. \tag{3}$$

The weights, w_m^+, are calculated as

$$w_m^+ = \frac{\phi\left(\frac{P_m^{orig} - 0.5}{h^{orig}}, \frac{P_m^{dest} - \overline{P}^{dest}}{h^{dest}}\right)}{\sum_{j:\, P_j^{orig} \geq .5, P_j^{dest} < 0.5} \phi\left(\frac{P_j^{orig} - 0.5}{h^{orig}}, \frac{P_j^{dest} - \overline{P}^{dest}}{h^{dest}}\right)},$$

where $\phi(\cdot)$ is the bivariate standard normal pdf and h^{orig} and h^{dest} are bandwidths. The weights, w_m^-, are defined similarly. This process is then repeated for a range of values for $\overline{P}^{dest}$ to get an estimate of the treatment effect, $\tau_1^{orig}(P_m^{dest})$, along the entire treatment cutoff. The estimators of $\tau_1^{dest}(P_m^{orig})$, $\tau_1^{orig}(P_m^{dest})$, and $\tau_2^{dest}(P_m^{dest})$ are defined similarly.

To simplify the choice of bandwidth in multiple dimensions, we transform the predictors of coverage prior to estimation to have mean zero and identity covariance matrix (see Pagan & Ullah, 1999). This allows us to check the sensitivity

[21] See Imbens and Lemieux (2008) for an introduction to RD and Hahn, Todd, and Van Der Klaauw (2001) for a detailed discussion of identification of treatment effects within an RD framework.

FIGURE 5.—SMOOTHED SURFACES, AVERAGE FARES, AND LCC PENETRATION

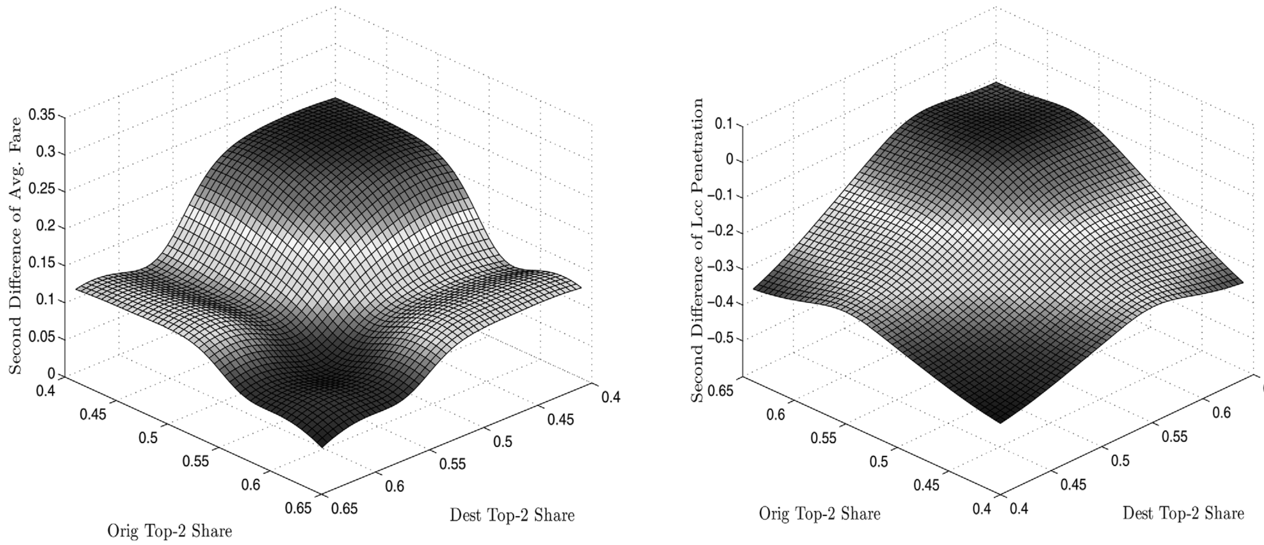

(a) Avg. Fare (b) Lcc Penetration

of our results by varying a single factor of proportionality, k, such that both h^{orig} and h^{dest} are equal to

$$h = kN^{-\frac{1}{4+d}},$$

where N is the number of observations in the quadrant of interest and $d = 2$ (dimension of predictor support). In equations (2) and (3),

$$N^{+} = \sum_{m} 1[P_m^{orig} < .5, P_m^{dest} < .5]$$

and

$$N^{-} = \sum_{m} 1[P_m^{orig} \geq .5, P_m^{dest} < .5],$$

respectively. We find our results to be insensitive to the choice of bandwidth and use $k = 2$ for all the RD analysis, which allows a great deal of flexibility, as we will discuss below, yet adequately smooths the surface.[22]

Calculating asymptotically valid standard errors for our estimates is a nontrivial computational exercise for a number of reasons. First, we are estimating a nonparametric surface in multiple dimensions. Second, we are most interested in the estimates of this nonparametric surface at the coverage cutoffs. Finally, we must account for the dependence in our data resulting from markets having end points in common, as was the case in the window regressions. For these reasons, we appeal to the resampling with dependent data literature to calculate asymptotically valid point-wise standard errors. (For a detailed treatment of resampling techniques for dependent data, see Lahiri, 2003). The clear dependence structure in our data makes application of these techniques straightforward. We treat the sample as representative of the population and compute jack-knife standard errors where we leave out blocks of markets with a common end point. Precisely, for each airport we find all markets with a common end point and drop them from the sample. Using the resulting subsample, we then reestimate the model. We repeat this process for each market and use the distribution of the estimates across subsamples to infer moments of the asymptotic distribution of our LATEs.

Results. The results and conclusions of our RD analysis are consistent and nearly identical to our findings using the window regression approach. For this reason, we focus the discussion of our RD results on the impact of coverage on fares and low-cost competition. Figures 5a and 5b present the surfaces relating second differences of the average fare and low-cost penetration, $\Delta_t Avg.Fare_{ijmt}$ and $\Delta_t Lcc.Penetration_{mt}$, respectively, to the predictors of treatment, where we smooth over the policy discontinuity. Foreshadowing the results from our RD analysis, we find the gradient of the surface changes rapidly around the treatment cutoffs in a way that is consistent with the results from the window regressions. In particular, there are four clearly visible planes in each of the surfaces, corresponding to the four quadrants in figure 1.

Figure 6 shows the estimated surface for a representative region of the surface relating the second difference of average fares to the predictors of treatment, along with the data used in estimation in the same region, when we allow for a break at the coverage cutoff. The figure makes clear that our RD is not as clear as many in the literature, in the sense that the

[22] We also explored cross-validation methods for choosing k, but we find it performs very poorly in our application by suggesting a bandwidth near 0 that overfits the data. A few aspects of our application and the method (interdependence of observations, multiple dimensional predictor vector, and the slow convergence rate of the cross-validation method) make this an unsurprising result.

FIGURE 6.—LOCAL COVERAGE EFFECT ON AVERAGE FARE

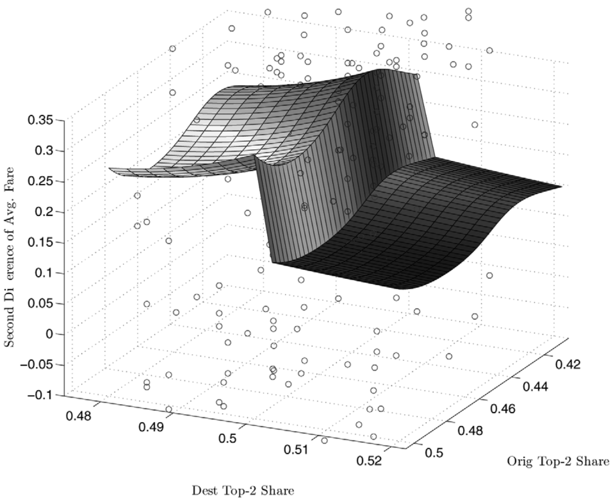

discontinuity is not plainly visible. This is not surprising due to the high variability in fares, two-dimensional support of our predictors, and the clumpiness of the data. This motivates our two-dimensional smoothing procedure.

Figures 7a and 7b and table 10 present the complete results of our RD analysis of fares and low-cost entry. Figures 7a and 7b show the RD surfaces, relating average fares and low-cost entry to the predictors of treatment, respectively, while allowing for a break in the surface at the coverage cutoffs. Not surprisingly, given the smoothed surfaces in Figures 5a and 5b and our findings from the window regressions, we find economically and statistically significant breaks along the coverage cutoffs for both market outcomes. Table 10 presents the point estimates along with standard errors, calculated using the block-bootstrapping procedure already discussed. Collectively, the results of our RD analysis both corroborate and provide additional insights to the findings gleaned from the window regressions.

The RD estimates of the effect on average fares of going from no end points treated to one end point treated range from 12% to 15%, only slightly larger than our findings from the window regressions. The effect of going from no coverage to both end points covered, if measured as the jump right at the vertex of the four surfaces, is about 27% (15.8 plus 11.6 or 15.6 plus 11.3). While, strictly speaking, this is the only correct way to measure this effect, averaging over the surfaces suggests the effect is around 18% to 21%. The reason for this difference is that unlike the window regressions, we don't impose symmetric effects for the one-end-point treated markets, and as figure 7a shows, we find substantial heterogeneity in the LATEs for the effect of coverage for the second market endpoint. Specifically, the effect of coverage is substantial and statistically significant for less concentrated airports (about 0.5), while not statistically different from 0 for more concentrated airports (0.6 or greater).[23]

For low-cost penetration, the magnitudes of the estimates are slightly smaller and exhibit some interesting patterns. One advantage of employing a true RD approach in our application is the opportunity to look for heterogeneity in the effect of coverage. This heterogeneity is obvious as one looks at the point estimates of the effect from coverage on low-cost penetration in table 10. This heterogeneity is also clearly visible in figure 7b. Looking at the estimates of $\tau_1^{orig}(P_m^{dest})$, we find a nearly monotonic relationship between concentration and the effect of coverage on low-cost entry, ranging from a point estimate of 0.024 to 0.218. As one would expect, coverage had its largest effect on entry in more concentrated markets. This explains why we get a statistically insignificant point estimate for the coefficient on $1[1\ \text{cover}_m]$ in the window regressions, where we can identify only an average effect of coverage. The effect of going from no end point to both end points being covered, again measured at the jump of the vertex between the four surfaces, is 0.28 and 0.31. This implies that approximately three additional markets out of every ten, with a two-firm concentration of above 0.5, have a low-cost presence as a result of AIR-21. Consistent with the fare estimates, we find only a moderate increase in low-cost penetration in the most concentrated markets when the second end point is covered. Collectively, our results suggest that coverage had the largest effect, fare declines and increased low-cost penetration, on moderately concentrated airports.

Regression-Discontinuity Validity. Above, we have discussed why we are comfortable assuming there are no (local) selection effects associated with AIR-21. The validity of our identifying assumption also requires there be no problem with incentive effects; that is, that carriers do not manipulate enplanement levels to avoid treatment. There are a number of reasons that we believe this is a valid assumption. First, coverage is determined at the airport level, not the airline level. Therefore, no individual airline can manipulate enplanements and entirely determine coverage; rather, it would take a cooperative effort by airlines serving the airport. Second, coverage in each year was determined using FAA enplanement data from two years earlier. An airline attempting to avoid coverage by the legislation would have been required to foresee the exact details of the legislation (including the exact enplanement cutoff) two years in advance of its passage.[24] Finally, manipulating enplanements at any one airport, particularly a large airport, has significant costs to an airline in terms of adjusting traffic in its entire network.

Extending formal tests to check for the strategic manipulation of enplanements (see McCrary, 2008) with a two-dimensional predictor vector is not immediately clear. However, we develop an informal test for manipulation of the predictors of treatment and provide evidence that little or no

[23] The symmetry of the surface remains, and the results are nearly identical if the return legs of round-trip tickets are not used in computing fares.

[24] We also examined enplanements of airports near the coverage cutoff in the years prior to AIR-21, 1998 to 2000. We find no evidence that concentration measures changed in a way that would suggest airports or carriers knew of the particular value of the coverage cutoff.

FIGURE 7.—RD SURFACE, AVERAGE FARE, AND LCC PENETRATION

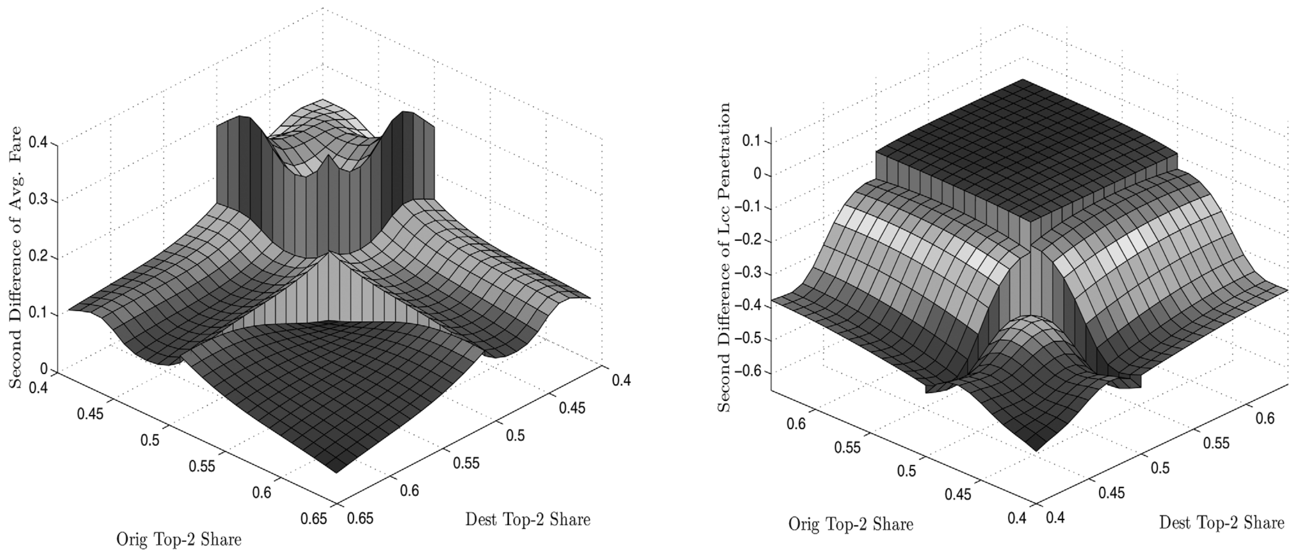

(a) Avg. Fare (b) Lcc Penetration

TABLE 10.—RDD COVERAGE ESTIMATES, AVG.FARE AND LCC.PRESENCE

Predictor	$\tau(P^{orig} < 0.5, P^{dest} = 0.5)$		Predictor	$\tau(P^{orig} < 0.5, P^{dest} = 0.5)$	
(P^{orig}, P^{dest})	Log(Avg.Fare)	Lcc.Presence	(P^{orig}, P^{dest})	Log(Avg.Fare)	Lcc.Presence
(0.4,0.50)	−0.128***	0.024	(0.50,0.40)	−0.136***	0.038
	(0.019)	(0.033)		(0.02)	(0.032)
(0.42,0.50)	−0.153***	0.017	(0.50,0.42)	−0.168***	0.048*
	(0.02)	(0.03)		(0.02)	(0.03)
(0.44,0.50)	−0.168***	0.046*	(0.50,0.44)	−0.177***	0.085***
	(0.021)	(0.026)		(0.019)	(0.027)
(0.46,0.50)	−0.153***	0.112***	(0.50,0.46)	−0.151***	0.145***
	(0.019)	(0.022)		(0.017)	(0.024)
(0.48,0.50)	−0.144***	0.158***	(0.50,0.48)	−0.143***	0.183***
	(0.017)	(0.023)		(0.016)	(0.026)
(0.50,0.50)	−0.17***	0.197***	(0.50,0.50)	−0.173***	0.218***
	(0.019)	(0.028)		(0.018)	(0.031)
(0.50,0.50)	−0.12***	0.07***	(0.50,0.50)	−0.124***	0.092***
	(0.014)	(0.024)		(0.014)	(0.021)
(0.52,0.50)	−0.09***	0.053***	(0.50,0.52)	−0.094***	0.081***
	(0.012)	(0.023)		(0.013)	(0.021)
(0.54,0.50)	−0.063***	0.041*	(0.50,0.54)	−0.066***	0.074***
	(0.012)	(0.023)		(0.013)	(0.022)
(0.56,0.50)	−0.039***	0.034	(0.50,0.56)	−0.043***	0.071***
	(0.011)	(0.023)		(0.012)	(0.022)
(0.58,0.50)	−0.021***	0.032	(0.50,0.58)	−0.024***	0.07***
	(0.011)	(0.023)		(0.012)	(0.022)
(0.60,0.50)	−0.007	0.034	(0.50,0.60)	0.011	0.072***
	(0.011)	(0.023)		(0.011)	(0.022)
(0.62,0.50)	0.004	0.039*	(0.50,0.62)	0.011	0.076***
	(0.011)	(0.023)		(0.011)	(0.022)
(0.64,0.5)	0.013	0.048***	(0.50,0.64)	0.011	0.082***
	(0.011)	(0.023)		(0.011)	(0.022)

Statistical significance: ***1%, **5%, *10%.

strategic manipulation of enplanements occurred. The test is based on the simple observation that airports just below the coverage cutoff in which one carrier controls a larger proportion of the traffic will be most vulnerable to strategic manipulation of enplanements. For example, consider two airports where the two largest carriers enplane 49% of the passengers. Suppose at the first airport, the top carrier enplanes 35% of all passengers, while the top carrier at the second airport enplanes 25% of all passengers. If an airline was attempting to avoid coverage of an airport by AIR-21 by manipulating enplanements, one would expect this to occur at the first airport. At the first airport, the largest carrier would have greater control in ensuring that the airport was not covered.

FIGURE 8.—TEST FOR STRATEGIC MANIPULATION

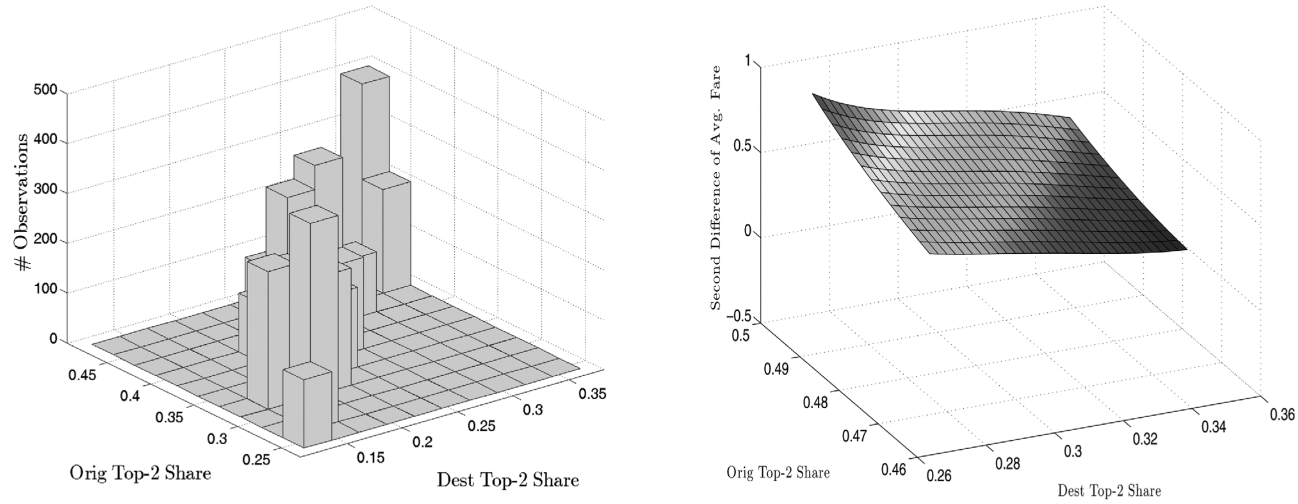

(a) Density of Observations (b) Avg. Fare

One way a carrier can lower enplanements is by raising fares. If a carrier was seeking to raise fares and lower its share of enplanements to avoid coverage, one would expect to see less of a drop in fares in markets near the coverage cutoff where one carrier has a larger share of enplanements. Figures 8a and 8b show no evidence to support a claim that enplanements were manipulated. In figure 8a, we plot the joint density of the share of the two largest carriers and the share of the largest carrier for airports below the cutoff. Given the high correlation between these two variables, we can plot only the relationship between these variables and changes in fares (difference between average fare in the first quarter of 2008 and 2001) for a small range of values in figure 8b. If carriers chose to strategically manipulate enplanements, we would expect to see a surface sloping up in the top carrier's share at airports closest to the cutoff. We find no evidence to support this claim.

V. Conclusion

High fares at concentrated airports have been a fact of life in U.S. air travel since the deregulation of the industry in 1979. The welfare implications of these high fares are ambiguous because consumers value both the size and scope, in the form of frequency and network size, of an airline when flying out of their home airport. However, size and scope lead to market power due to scarce airport facilities. In 2000, the U.S. Congress took a stand on the question, deciding too much market power at highly concentrated airports was generating too much of the fare difference and enacted AIR-21. Among other things, AIR-21 required concentrated airports to take steps to increase competition and make airport facilities available to all carriers wanting to serve the airport.

We have provided evidence that AIR-21 was successful in encouraging new and intensified competition at its targeted airports. Moreover, we have found evidence that Congress was right in concluding that market power contributed too much to high fares from the perspective of consumers. That is, we find little evidence that increased competition has significantly eroded quality provision, either directly by reducing large incumbent size or indirectly by disincentiving high frequencies. The only meaningful degradation of service quality we find has come in the on-time performance of airlines at treated airports. The magnitude of these changes in service quality relative to the magnitude of fare declines, along with the fact that we see increases in passenger volumes, suggest AIR-21 had a strongly positive effect on consumer welfare.

Our quasi-experimental approach to analyzing the impact of barriers to entry is somewhat novel in the industrial organization literature, and we think our clean identification strategy represents a significant contribution to it. However, our study also highlights some of the difficulties in implementing such a research design. While we can explain between 40% and 50% of the decline in fares in covered markets, a result of intensified competition from low-cost carriers, it remains an open question to identify other determinants. Moreover, if we had arrived at more nuanced results (e.g., more significant declines in quality and no increase in passenger volumes), we would need more structure to say much about the balance of welfare gains and losses.

The competition plans and subsequent FAA reports provide at least a subset of the actions taken by airports and seem to provide a good source for identifying other possible explanations. A couple of candidates that seem likely to have some explanatory power are the reduction of landing fees for smaller carriers to the levels enjoyed by large-presence incumbents as well as limits on subleasing fees that can be charged by one carrier to another for the use of underused boarding gates. Both steps, discussed in the majority of the airports' competition plans, could be a significant source of cost pass-throughs from carriers to consumers. Carriers may also simply reduce fares to generate outcomes consistent with the goals of AIR-21 to avoid additional oversight in the future. We leave more detailed investigation of these channels for future research.

REFERENCES

Berry, Steven, "Airport Presence as Product Differentiation," *American Economic Review P&P* 80 (1990), 394–399.

——— "Estimation of a Model of Entry in the Airline Industry," *Econometrica* 60 (1992), 889–917.

Berry, Steven, Michael Carnall, and Pablo Spiller, "Airline Hubs: Costs, Markups, and the Implications of Costumer Heterogeneity," in Darin Lee, ed., *Advances in Airline Economics, Competition Policy and Antitrust*, vol. 1 (New York: Elsevier, 2006).

Berry, Steven, and Panle Jia, "Tracing the Woes: An Empirical Analysis of the Airline Industry," *American Economic Journal: Microeconomics* 2 (2010), 1–43.

Black, Sandra, "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics* 114 (1999), 577–599.

Borenstein, Severin, "Hubs and High Fares: Dominance and Market Power in the U.S. Airline Industry," *Rand Journal of Economics* 20 (1989), 344–365.

Borenstein, Severin, and Nancy Rose, "Competition and Price Dispersion in the U.S. Airline Industry," *Journal of Political Economy* 102 (1994), 653–683.

Bresnahan, Timothy, and Peter Reiss, "Entry and Competition in Concentrated Markets," *Journal of Political Economy* 99 (1991), 977–1009.

Brueckner, Jan, Darin Lee, and Ethan Singer, "Airline Competition and Domestic U.S. Airfares: A Comprehensive Reappraisal," University of California, Irvine, working paper (2011).

Brueckner, Jan, and Pablo Spiller, "Economies of Traffic Density in the Deregulated Airline Industry," *Journal of Law and Economics* 37 (1994), 379–415.

Cameron, A. Colin, Jonah Gelbach, and Douglas Miller, "Robust Inference with Multi-Way Clustering," NBER technical working paper T0327 (2007).

Ciliberto, Federico, and Elie Tamer, "Market Structure and Multiple Equilibria in the Airline Industry," *Econometrica* 77 (2009), 1791–1828.

Ciliberto, Federico, and Jonathan Williams, "Limited Access to Airport Facilities and Market Power in the Airline Industry," *Journal of Law and Economics* 53 (2010), 467–495.

Evans, William, and Ioannis Kessides, "Living by the 'Golden Rule': Multimarket Contact in the U.S. Airline Industry," *Quarterly Journal of Economics* 109 (1994), 341–366.

Fan, Jianqing, and Irène Gijbels, *Local Polynomial Modelling and Its Applications* (London: Chapman & Hall, 1996).

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

Goolsbee, Austan, and Chad Syverson, "How Do Incumbents Respond to the Threat of Entry? Evidence from the Major Airlines," *Quarterly Journal of Economics* 123 (2008), 1611–1633.

Hahn, Jinyong, Petra Todd, and Wilbert Van Der Klaauw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica* 69 (2001), 201–209.

Halvorsen, Robert, and Raymond Palmquist, "The Interpretation of Dummy Variables in Semilogarithmic Equations," *American Economic Review* 70 (1980), 474–475.

Imbens, Guido, and Thomas Lemieux, "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics* 142 (2008), 615–635.

Lahiri, Soumendra, *Resampling Methods for Dependent Data* (New York: Springer, 2003).

Lederman, Mara, "Do Enhancements to Loyalty Programs Affect Demand? The Impact of International Frequent Flyer Partnerships on Domestic Airline Demand," *Rand Journal of Economics* 38 (2007), 1134–1158.

——— "Are Frequent Flyer Programs a Cause of the 'Hub Premium'?" *Journal of Economics and Management Strategy* 17 (2008), 35–66.

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

Mazzeo, Michael, "Product Choice and Oligopoly Market Structure," *Rand Journal of Economics* 33 (2002), 221–242.

McCrary, Justin, "Testing for Manipulation of the Running Variable in the Regression Discontinuity Design," *Journal of Econometrics* 142 (2008), 698–714.

Pagan, Adrian, and Aman Ullah, *Nonparametric Econometrics* (Cambridge: Cambridge University Press, 1999).

Seim, Katja, "An Empirical Model of Firm Entry with Endogenous Product-Type Choices," *Rand Journal of Economics* 37 (2006), 619–640.

Transportation Research Board, *Entry and Competition in the U.S. Airline Industry: Issues and Opportunities* (Washington, DC: Transportation Research Board, 1999).

U.S. General Accounting Office, *Barriers to Competition in the Airline Industry* (Washington, DC: U.S. General Accounting Office, 1989).

——— *Airline Competition: Industry Operating and Marketing Practices Limit Market Entry* (Washington, DC: U.S. General Accounting Office, 1990).

——— *Challenges in Enhancing Competition in Dominated Markets* (Washington, DC: U.S. General Accounting Office, 2001).

Williams, Jonathan, "Capacity Investments, Exclusionary Behavior, and Welfare: A Dynamic Model of Competition in the Airline Industry," *SSRN* working paper 1375450 (2012).