The Role of Route-Level Decisions in the Efficiency and Resilience of Airline Operations: Evidence from the Wright Amendment Repeal

Vishal Ahuja, Yasin Alan, Mazhar Arıkan

A. Supplementary Material

A.1. Overview of Publicly Available Datasets

Two datasets published by the Bureau of Transportation Statistics (BTS) are relevant to our dataset. In this section, we first describe these two datasets and compare them with our dataset. We then describe two additional datasets we use to construct our explanatory and control variables.

Airline On-Time Performance (OTP) Data: This dataset contains flight-level information reported by U.S. air carriers that exceed a revenue threshold. Until 2018, that threshold was 1% of domestic scheduled passenger revenues, but it changed to 0.5% in 2018. BTS publishes OTP data monthly. Each flight entry (row) in the dataset includes more than 100 columns, including the flight date, carrier information, airplane tail number, information about origin and destination airports, scheduled and actual departure and arrival times, binary codes for canceled and diverted flights, causes of delay and cancellation, and non-stop distance. For example, Table A1 provides selected flight-level information for the three flights we use in our itinerary-level examples Table 1 in the main body. While we use gate departure and arrival times in our analysis, Table A1 provides additional information about taxi times, wheels off time (i.e., when the aircraft leaves the ground at the origin airport), and wheels on time (i.e., when the aircraft's wheels touch the ground at the destination airport) along with aircraft tail numbers and non-stop distances between origin and destination airports. More information about OTP data is available at the BTS website, `https://www.transtats.bts.gov/tables.asp?qo_vq=EFD&QO_anzr=`.

As explained in Section 1.1, OTP data do not include passenger-level details, limiting their use for capturing passengers' travel experiences. For example, OTP data do not report the number of passengers who encountered a 36-minute arrival delay in their non-stop flight from DAL to AUS. By contrast, our dataset reveals that 43 passengers experienced this delay. Similarly, the OTP data would not be sufficient to conclude that the third passenger in Table 1 in the main body missed their connection in AUS.

Airline Origin and Destination Survey (DB1B): DB1B is a 10% sample of airline tickets from reporting carriers. BTS collects and publishes these data quarterly. Each DB1B ticket entry includes information on the origin, destination, and layover airports, as well as the number of passengers, operating carriers, fares, and travel distances. For example, the DB1B data for the second quarter of 2014, which covers the flight date of the itinerary-level examples we provide in

Table A1 OTP Data Examples

	DAL-AUS	AUS-LAS	AUS-DEN
FlightDate	2014-05-30	2014-05-30	2014-05-30
Carrier	WN	WN	WN
TailNum	N299WN	N487WN	N943WN
Origin	DAL	AUS	AUS
Dest	AUS	LAS	DEN
CRSDepTime	1750	2025	1920
DepTime	1824	2021	1918
TaxiOut	12	13	23
WheelsOff	1836	2034	1941
WheelsOn	1916	2113	2028
TaxiIn	5	6	7
CRSArrTime	1845	2120	2035
ArrTime	1921	2119	2035
ArrDelay	36	-1	0
CRSElapsedTime	55	175	135
ActualElapsedTime	57	178	137
Distance	189	1090	775

Notes. FlightDate is the date of flight, Carrier is a code that identifies the airline (Southwest's carrier code is WN), TailNum is the tail number of the airplane, Origin is the airport code for the origin airport, Dest is the airport code for the destination airport, CRSDepTime is the scheduled gate departure time (local time: hhmm), DepTime is the actual gate departure time (local time: hhmm), TaxiOut is the taxi out time (in minutes), WheelsOff is wheels off time (local time: hhmm), WheelsOn is wheels on time (local time: hhmm), TaxiIn is the taxi in time (in minutes), CRSArrTime is the scheduled gate arrival time (local time: hhmm), ArrTime is the actual gate arrival time (local time: hhmm), ArrDelay is the difference between ArrTime and CRSArrTime (in minutes), CRSElapsedTime is the scheduled flight time (in minutes), ActualElapsedTime is the actual flight time (in minutes), and Distance is the distance between the origin and destination airports (in miles).

Table 1 in the main body (May 30, 2014), include ticket information for 17 Southwest passengers who traveled from DAL to DEN with a layover in AUS. Table A2 presents five ticket examples from DB1B. Revisiting the itinerary examples in Table 1 in the main body, it is unclear if the third passenger's ticket is in DB1B, which constitutes a 10% sample of airline tickets. By contrast, our dataset reveals that 1,931 Southwest passengers traveled from DAL to DEN with a layover in AUS in the second quarter of 2014. Notably, our dataset provides additional itinerary-level details, including travel dates and scheduled flight and layover times, which DB1B lacks. More information about DB1B is available at the BTS website, https://www.transtats.bts.gov/DatabaseInfo.asp?QO_VQ=EFI&Yv0x=D

Table A2 DB1B Data Examples

Year	Quarter	AirportGroup	OpCarrierGroup	Passengers	MktFare	MktMilesFlown	NonStopMiles
2014	2	DAL:AUS:DEN	WN:WN	1	241	964	651
2014	2	DAL:AUS:DEN	WN:WN	1	357	964	651
2014	2	DAL:AUS:DEN	WN:WN	1	122.37	964	651
2014	2	DAL:AUS:DEN	WN:WN	2	418.43	964	651
2014	2	DAL:AUS:DEN	WN:WN	1	263.89	964	651

Notes. AirportGroup is the passenger's route (e.g., DAL:AUS:DEN is a trip from DAL to DEN with a layover in AUS), OpCarrierGroup is the list of carriers operating each leg (e.g., WN:WN means Southwest operated both legs), Passengers is the number of passengers on a ticket, MktFare is the fare, MktMilesFlown is the total flight distance (in miles) and NonStopMiles is the non-stop distance between origin and destination airports (in miles).

Weather Data: We construct our weather control variables using the Iowa State University

Automated Airport Weather Observations Database. We use the `riem` library in `R` to download weather variables from this database. Specifically, the `riem_measures` function takes the airport code and the start and end dates as inputs and returns weather variables, including air temperature, wind speed, precipitation, and visibility. More information about this database is available at the `riem` library's vignette, https://cran.r-project.org/web/packages/riem/riem.pdf.

FAA Operations and Performance (OP) Data: These data contain historical traffic counts, forecasts of aviation activity, and delay statistics. Specifically, we use the Aviation System Performance Metrics (ASPM) dataset, which includes data on flights to and from 77 major airports tracked in the ASPM dataset. The variable we use for the arrival (departure) capacity is the airport arrival rate (AAR; airport departure rate, ADR). AAR (ADR) is the airport-determined arrival (departure) rate the airport can handle per unit of time based on current weather conditions, traffic mix, and runway configuration. Using this dataset, we obtain the AAR and ADR of each airport for each month between 2013 and 2015. From these records, we calculate each airport's average hourly arrival (departure) capacities. More information about the ASPM database is available at the FAA website, https://aspm.faa.gov/.

A.2. Southwest's Network During The Study Period

During our study period between October 13, 2013, and October 12, 2015, Southwest's network had 20 type I, 33 type II, and 35 type III destinations. The type I destinations are Albuquerque, NM (ABQ); Amarillo, TX (AMA); Austin, TX (AUS); Birmingham, AL (BHM); Branson, MO (BKG); Corpus Christi, TX (CRP); El Paso, TX (ELP); Harlingen/San Benito, TX (HRL); Houston, TX (HOU); Jackson/Vicksburg, MS (JAN); Kansas City, MO (MCI); Little Rock, AR (LIT); Lubbock, TX (LBB); Midland/Odessa, TX (MAF); New Orleans, LA (MSY); Oklahoma City, OK (OKC); San Antonio, TX (SAT); St. Louis, MO (STL); Tulsa, OK (TUL); and Wichita, KS (ICT).

The type II destinations are Akron, OH (CAK); Albany, NY (ALB); Boise, ID (BOI); Buffalo, NY (BUF); Burbank, CA (BUR); Cleveland, OH (CLE); Dayton, OH (DAY); Des Moines, IA (DSM); Flint, MI (FNT); Fort Myers, FL (RSW); Grand Rapids, MI (GRR); Greer, SC (GSP); Hartford, CT (BDL); Islip, NY (ISP); Jacksonville, FL (JAX); Key West, FL (EYW); Louisville, KY (SDF); Manchester, NH (MHT); Minneapolis, MN (MSP); Newark, NJ (EWR); Norfolk, VA (ORF); Ontario, CA (ONT); Pensacola, FL (PNS); Portland, ME (PWM); Providence, RI (PVD); Reno, NV (RNO); Richmond, VA (RIC); Rochester, NY (ROC); San Juan, PR (SJU); Spokane, WA (GEG); Tucson, AZ (TUS); Washington, DC (IAD); and West Palm Beach/Palm Beach, FL (PBI).

The type III destinations are Atlanta, GA (ATL); Baltimore, MD (BWI); Boston, MA (BOS); Charleston, SC (CHS); Charlotte, NC (CLT); Chicago, IL (MDW); Columbus, OH (CMH); Denver, CO (DEN); Detroit, MI (DTW); Fort Lauderdale, FL (FLL); Indianapolis, IN (IND); Las Vegas, NV (LAS); Los Angeles, CA (LAX); Memphis, TN (MEM); Milwaukee, WI (MKE); Nashville, TN (BNA); New York, NY (LGA); Oakland, CA (OAK); Omaha, NE (OMA); Orlando, FL (MCO); Panama City, FL (ECP); Philadelphia, PA (PHL); Phoenix, AZ (PHX); Pittsburgh, PA (PIT); Portland, OR (PDX); Raleigh/Durham, NC (RDU); Sacramento, CA (SMF); Salt Lake City, UT (SLC); San Diego, CA (SAN); San Francisco, CA (SFO); San Jose, CA (SJC); Santa Ana, CA (SNA); Seattle, WA (SEA); Tampa, FL (TPA); and Washington, DC (DCA).

A.3. Displacement Procedure

In this section, we provide a detailed explanation of our displacement procedure for connecting passengers. Our procedure has three steps. In the first step, we follow the approach of Herring et al. (2019) to determine airport-specific minimum layover times. These layover times serve as a proxy for the time a passenger needs to catch their connecting flight. For example, analyzing all itineraries with layovers at LGA reveals that the minimum scheduled layover time at LGA was 30 minutes (i.e., $MSLT_{LGA} = 30$). However, after conducting a similar analysis for all airports in Southwest's network, we found cases in which the minimum scheduled layover time was notably short (e.g., as little as five minutes). Considering practical factors, such as the time it takes to deplane after an aircraft arrives at the gate and the boarding process ending before an aircraft's gate departure, we recognize that passengers are unlikely to make their connection if the difference between the actual gate departure time of the second leg and the actual gate arrival time of the first leg is less than 20 minutes. This holds regardless of an airport's size or congestion level. Therefore, if an airport's minimum scheduled layover time is below 20 minutes in our dataset, we adjust it to 20 minutes.

After calculating each airport's minimum scheduled layover time, we proceed to the second step. In this step, we identify the disrupted itineraries by comparing each connecting passenger's actual layover time with the minimum scheduled layover time of the connection airport. Specifically, a passenger's itinerary is considered disrupted at airport a if their actual layover time at that airport (i.e., the difference between the actual gate departure time of the flight departing from airport a and the actual gate arrival time of the flight arriving at airport a) is less than the minimum scheduled layover time of that airport ($MSLT_a$). For example, revisiting Table 1 in the main body, the actual layover time of the third passenger is -3 minutes (since the actual arrival time of the first leg is 19:21 and the actual departure time of the second leg is 19:18), which is less than $MSLT_{AUS} = 20$ minutes. Consequently, this passenger's itinerary is disrupted.

After identifying all disrupted passengers, we move to the third step of our displacement procedure. In this step, we reallocate each disrupted passenger to the next available flight scheduled to arrive at their final destination. For example, the third passenger in Table 1 in the main body missed their flight from AUS to DEN on May 30, 2014. Southwest's subsequent flight from AUS to DEN was scheduled for the next day (May 31, 2014) at 8:05. We rebooked the passenger on that flight, which arrived at DEN at 9:03. Comparing this to the passenger's original itinerary, where the scheduled arrival time at DEN was 20:35 on May 30, 2014, we observe a significant delay. Specifically, the actual arrival time of 9:03 on May 31, 2014, resulted in a 12-hour and 28-minute delay. Since this delay exceeds three hours, our binary delay variable, *Delayed*, is set to 1 for this passenger.

Admittedly, Southwest may have provided an alternative accommodation to that passenger (e.g., rerouting the passenger from AUS to another airport and then to DEN or placing them on a different flight from AUS to DEN due to seat unavailability on the 8:05 flight). While alternative accommodations can impact the length of the delay, they are unlikely to significantly affect our main results. This is because our resilience analysis relies on binary classification, where $Delayed_i = 1$ if passenger i's delay exceeds three hours. Our conversations with airline managers have revealed that it is rare for a disrupted passenger to reach their final destination within three hours of their originally scheduled arrival time. In other words, when an airline rebooks a passenger due to a missed connection, $Delayed_i$ is likely to be 1, regardless of the details of the rebooking process.

A.4. Variable Construction and Additional Summary Statistics

In this section, we first provide detailed descriptions of the key variables used in our empirical models and present some numerical examples. Then, we provide summary statistics for the variables we construct in Sections 3 and 4. Table A3 provides detailed descriptions of the key variables used in our empirical models along with some numerical examples.

Table A3: Variable Definitions

Variable	Definition
L_i	Number of legs in passenger i's itinerary. Revisiting Table 1 in the main body, L_i equals one for the first passenger and two for the second and third passengers. As a side note, in some cases, a passenger takes a "direct" flight with the plane making a scheduled stop en route to the passenger's final destination. In such cases, the scheduled stop counts as a layover even though the passenger does not change planes during the stop.
SDT_{il} (SAT_{il})	Scheduled departure (arrival) time of leg $l = 1, \cdots, L_i$ in passenger i's itinerary.
AAT_{il}	Actual arrival time of $l = 1, \cdots, L_i$ in passenger i's itinerary.
STT_i	Scheduled travel time for passenger i. After adjusting for the time zone difference, we calculate the scheduled travel time for passenger i as $STT_i = SAT_{i,L_i} - SDT_{i,1}$. For example, the scheduled travel time for the second passenger in Table 1 in the main body is 5.5 hours.
$Delayed_i$	Binary variable that equals one if passenger i's travel delay (i.e., the difference between the actual and scheduled arrival times of the last leg) exceeds three hours. When a passenger is displaced, we calculate this variable after rebooking the passenger based on our displacement procedure. For example, as we explain in greater detail in Section A.3, $Delayed_i = 1$ for the third passenger in Table 1 in the main body.
$Dallas_i$	Binary variable that equals one if passenger i's origin was DAL.
$Post_i$	Binary variable that equals one if passenger i's itinerary date was after the repeal.
$Treated_i$	Interaction effect $Dallas_i \times Post_i$.

Continued on next page

Table A3: Variable Definitions (Continued)

Variable	Definition
$distance_{il}$	Non-stop distance between the origin and destination airports of leg l in passenger i's itinerary. For example, for the third passenger in Table 1 in the main body, $distance_{il}$ is 189 miles for the first leg (from DAL to AUS) and 775 miles for the second leg (from AUS to DEN), as the OTP records in Table A1 show.
$non.stop.distance_i$	Non-stop distance between the origin and destination airports for passenger i. For example, for the third passenger in Table 1 in the main body, the non-stop distance from origin (AUS) to destination (DEN) is 651 miles, as the DB1B records in Table A2 show.
Min. temperature	Minimum temperature on the date of passenger i's trip in the airports the passenger visits. For example, suppose that a passenger travels from DAL to DEN through a layover in AUS and that the minimum temperatures at DAL, AUS, and DEN on the day of that passenger's trip are 55, 45, and 30° Fahrenheit, respectively. Then, the temperature control variable for that passenger equals $\min\{55, 45, 30\} = 30$.
Max. wind	Maximum wind on the date of passenger i's trip in the airports the passenger visits.
Max. precipitation	Minimum precipitation on the date of passenger i's trip in the airports the passenger visits.
Min. visibility	Minimum visibility on the date of passenger i's trip in the airports the passenger visits.
$MSLT_a$	Minimum scheduled layover time for airport a. As we explain in greater detail in Section A.3, we first calculate the minimum scheduled layover time for airport a based on all itineraries with a layover at that airport. However, when the value we obtain is less than 20 minutes, we set it to 20 minutes to rule out impractical layover times.
DC_{il} (AC_{il})	Departure (arrival) congestion of leg l in passenger i's itinerary. To calculate DC_{il} (AC_{il}), we first count the number of flights scheduled to depart (arrive) between 45 minutes before and 15 minutes after the scheduled departure (arrival) time of that flight using OTP data. We then obtain the average hourly departure (arrival) capacity of the origin airport from the FAA's Operations and Performance Data. Lastly, we calculate DC_{il} (AC_{il}) as the ratio of the number of departures (arrivals) we count within the one-hour time block to the average hourly departure (arrival) capacity of the departure (arrival) airport. Because DC_{il} and AC_{il} capture airport utilization, they should vary between zero and one. In some rare instances (less than 1% of our observations), our congestion variables exceed one because the daily capacity of an airport is set very low by the FAA. In such cases, we set the congestion variables equal to one. For example, consider the flight from DAL to AUS in Table 1 in the main body. The flight was scheduled to depart from DAL at 17:50 and arrive at AUS at 18:45 on May 30, 2014. During the period from 45 minutes before to 15 minutes after the scheduled departure (arrival) time, there were a total of 10 (10) flights scheduled to depart from (arrive at) DAL (AUS). The average hourly departure capacity of DAL was 26.5 flights. Therefore, the departure congestion for this flight (DC_{il}) is calculated as the ratio of scheduled departures to the departure capacity, resulting in a congestion value of 10/26.5=0.377. The average hourly arrival capacity at AUS was 46 flights. Thus, the arrival congestion for this flight (AC_{il}) is calculated as the ratio of scheduled arrivals to the arrival capacity, yielding a congestion value of 10/46=0.217.
Max. congestion	Maximum arrival or departure congestion passenger i experiences during their trip, $\max\{AC_{i1}, \cdots, AC_{i,L_i}, DC_{i1}, \cdots, DC_{i,L_i}\}$.
$flight_{il}$	Number of non-stop flights from the origin to the destination of leg l on passenger i's travel date. For example, for the second passenger in Table 1, who traveled from DAL to LAS with a layover in AUS, this variable equals 12 for the first leg and three for the second leg. That is, according to OTP data, Southwest had 12 non-stop flights from DAL to AUS and three non-stop flights from AUS to LAS on May 30, 2014.

Continued on next page

Table A3: Variable Definitions (Continued)

Variable	Definition
$NomTT_i$	Nominal travel time is the sum of nominal flight and layover times for passenger i. To calculate the nominal flight time between two airports, we use the 10th percentile values of the gate-to-gate scheduled flight times between these airports. For example, consider the itinerary record of the second passenger in Table 1 in the main body, who traveled from DAL to LAS with a layover in AUS. The 10th percentile of gate-to-gate flight times for all flights between DAL and AUS is 45 minutes. Similarly, the 10th percentile of gate-to-gate flight times for all flights between AUS and LAS is 162 minutes. To calculate the nominal layover time at an airport, we analyze all itineraries with a layover at that airport and use the 10th percentile of the scheduled layover times. For each itinerary, we calculate the scheduled layover time by subtracting the first leg's scheduled arrival time from the second leg's scheduled departure time. For example, the second passenger in Table 1 in the main body had a 100-minute layover in AUS, as the scheduled arrival time of the first leg was 18:45 and the scheduled departure time of the second leg was 20:25. We extracted the scheduled layover times for each layover airport from all itineraries. For instance, the analysis of the itineraries using AUS as the layover airport reveals that the average scheduled layover time at AUS is approximately 69 minutes. Furthermore, the 10th percentile of layover times for itineraries using AUS as a layover airport is 25 minutes. Consequently, the nominal travel time for the second passenger in Table 1 in the main body is 45+25+162=232 minutes, which is the sum of the nominal travel time of the first leg (45), the nominal layover time (25), and the nominal travel time of the second leg (162).

Table A4 reports summary statistics for the variables we construct in Sections 3 and 4. While the summary statistics for $\overline{RI}_{o,d,t}$ show major changes for DAL-origin passengers, they are relatively stable for HOU-origin passengers. For example, the median $\overline{RI}_{o,d,t}$ for DAL-origin passengers is 2.10 before the repeal and 1.79 after. By contrast, the median $\overline{RI}_{o,d,t}$ for HOU-origin passengers is 1.6 before the repeal and 1.65 after. Indeed, the summary statistics for $ACRI_d$ and $ACRR_d$ indicate that DAL-origin passengers had access to more efficient and resilient routes to many destinations after the repeal.

B. Robustness Checks

B.1. Potential Violations of SUTVA

Our model specifications rely on SUTVA, which requires stable treatments (i.e., treatment is equally applied to treated passengers). Our study design ensures this because no passengers who departed from DAL after October 12, 2014, were subject to the Wright Amendment. In addition, SUTVA requires that there are no spillovers between the treatment and control groups. In our context, this assumption means that the repeal should not affect HOU-origin passengers. Although HOU-origin passengers were not subject to the Wright Amendment, its repeal might have an indirect impact on HOU-origin passengers by changing Southwest's propensity to use HOU (DAL) as a layover airport for its DAL-origin (HOU-origin) passengers.

We find that the repeal decreased the proportion of DAL-origin passengers using HOU as a layover airport from 13% to 5.1%, while the proportion of HOU-origin passengers using DAL as

Table A4 Additional Summary Statistics

	DAL pre-repeal							DAL post-repeal						
Metric	Mean	SD	P10	P25	Median	P75	P90	Mean	SD	P10	P25	Median	P75	P90
RI_i	1.59	0.55	1.13	1.19	1.36	1.86	2.36	1.50	0.45	1.19	1.25	1.36	1.53	2.05
RR_i	0.86	0.20	0.57	0.72	0.91	1	1.07	0.87	0.15	0.64	0.80	0.90	0.98	1.03
$\overline{RI}_{o,d,t}$	2.11	0.68	1.21	1.70	2.10	2.39	2.83	1.86	0.53	1.27	1.48	1.79	2.09	2.37
$\overline{RR}_{o,d,t}$	0.70	0.17	0.48	0.57	0.67	0.80	0.96	0.77	0.12	0.62	0.67	0.75	0.86	0.94

	HOU pre-repeal							HOU post-repeal						
Metric	Mean	SD	P10	P25	Median	P75	P90	Mean	SD	P10	P25	Median	P75	P90
RI_i	1.50	0.41	1.20	1.28	1.38	1.55	1.94	1.53	0.44	1.20	1.28	1.38	1.58	2.08
RR_i	0.84	0.15	0.62	0.76	0.87	0.93	1.01	0.85	0.15	0.64	0.76	0.88	0.94	1.01
$\overline{RI}_{o,d,t}$	1.80	0.51	1.30	1.40	1.60	2.04	2.49	1.81	0.48	1.31	1.45	1.65	2.08	2.45
$\overline{RR}_{o,d,t}$	0.76	0.13	0.58	0.66	0.77	0.85	0.92	0.78	0.12	0.63	0.69	0.78	0.86	0.95

	Destination-Level Metrics						
Metric	Mean	SD	P10	P25	Median	P75	P90
$ACRI_d$	−0.26	0.32	−0.69	−0.45	−0.25	−0.01	0.07
$ACRR_d$	0.06	0.11	−0.06	−0.01	0.04	0.14	0.20

Notes. The summary statistics for RI_i and RR_i are based on all passengers in the subsample. For example, the mean RR_i value for DAL-origin passengers during the pre-repeal period is 1.59. The summary statistics for $\overline{RI}_{o,d,t}$ and $\overline{RR}_{o,d,t}$ are based on destination-level averages. For example, for a given origin (HOU or DAL) and period (pre-repeal or post-repeal) pair, we first calculate $\overline{RI}_{o,d,t}$, which leads to one data point (the average RI score) for each destination. We then use the distribution of $\overline{RI}_{o,d,t}$ values to calculate the summary statistics we report in the table. The summary statistics for $ACRI_d$ and $ACRR_d$ are also based on one data point from each destination (i.e., the adjusted change in the average route inefficiency or route resilience for that destination).

a layover airport slightly increased from 4.7% to 5.6%. Because these proportions vary across destinations, we analyze passenger volumes at the destination level to identify the destinations that are most susceptible to a potential violation of SUTVA. Specifically, for each destination, we tabulate the number and proportion of DAL-origin passengers who use HOU as a layover airport during the pre- and post-repeal periods. For comparison purposes, we repeat the same calculations for HOU-origin passengers using DAL as a layover airport.

Formally, let $Passengers_{o-l-d,pre}$ and $Passengers_{o-l-d,post}$ denote the number of passengers traveling from origin $o \in \{DAL, HOU\}$ to destination d with a layover in $l \in \{DAL, HOU\}$, $l \neq o$ during the pre- and post-repeal periods, respectively. Then, we define the adjusted change in the connecting passenger volume to destination d as $ACCPV_d \equiv (Passengers_{DAL-HOU-d,post} - Passengers_{DAL-HOU-d,pre}) - (Passengers_{HOU-DAL-d,post} - Passengers_{HOU-DAL-d,pre})$. Table A5 shows the top 15 destinations with the largest adjusted decreases in the connecting passenger volume. For example, the number of DAL-origin passengers using HOU as a layover airport to travel to ATL sharply decreased from 31,000 (65% of all passengers traveling from DAL to ATL during the pre-repeal period) to 10,284 (8.36% of all passengers traveling from DAL to ATL during the post-repeal period), whereas the number of HOU-origin passengers using DAL as a layover airport to travel to ATL increased from 7 to 2,031. Note that our $ACCPV_d$ metric captures the changes for both DAL- and HOU-origin passengers. For example, LAX is the fifth destination on our list partly because a relatively large number of HOU-origin passengers traveled to LAX using DAL as a layover airport during the post-repeal period.

Table A5 Passenger Flows

		DAL-origin Passengers with a Layover in HOU				HOU-origin Passengers with a Layover in DAL				
		Pre-repeal		Post-repeal		Pre-repeal		Post-repeal		
	Destination	Count	Proportion (%)	Count	Proportion (%)	Count	Proportion (%)	Count	Proportion (%)	ACPPV
1	ATL	31,000	64.75	10,284	8.36	7	0.01	2,031	1.92	-22,740
2	LGA	18,716	57.15	2,195	1.87	6	0.01	3,753	3.67	-20,268
3	MCO	19,613	28.98	4,212	3.53	5	0.01	1,380	1.39	-16,776
4	TPA	18,522	42.57	6,258	8.68	1	0.00	1,984	3.27	-14,247
5	LAX	10,383	14.15	1,302	0.94	9	0.01	4,438	2.84	-13,510
6	BWI	12,906	17.88	2,387	2.27	6	0.01	1,620	2.04	-12,133
7	BNA	21,223	40.62	10,331	11.79	2	0.00	883	1.19	-11,773
8	FLL	11,548	40.68	4,079	5.12	0	0	3,757	6.88	-11,226
9	DCA	5,897	33.87	1,514	1.07	1	0.00	4,826	5.41	-9,208
10	MDW	8,604	7.74	1,184	0.52	28	0.02	1,738	0.91	-9,130
11	LAS	5,408	5.52	966	0.54	22	0.02	3,256	2.31	-7,676
12	EWR	11,277	48.20	5,859	34.30	6	0.01	17	0.02	-5,429
13	PHX	5,697	5.78	1,702	1.11	27	0.04	1,407	1.65	-5,375
14	SNA	3,785	40.69	1,309	2.96	0	0	2,360	6.53	-4,836
15	SAN	5,528	9.72	3,446	3.74	4	0.01	1,841	3.13	-3,919

Reestimating regression models on a subsample is commonly used in the literature to handle potential violations of SUTVA (e.g., Jo et al. 2020, Gillingham and Bollinger 2021). Accordingly, we rely on Table A5 and reestimate Models 1–6 on two subsamples. The first subsample excludes the top five destinations in Table A5 (i.e., ATL, LGA, MCO, TPA, and LAX) from our analysis. Table A6 reports the treatment effect coefficients for the first subsample. The second subsample excludes the top ten destinations in Table A5 (i.e., ATL, LGA, MCO, TPA, LAX, BWI, BNA, FLL, DCA, and MDW). Table A7 reports the treatment effect coefficients for the second subsample. The treatment effect coefficients in Tables A6 and A7 are in line with their counterparts in Tables 4 and 5. Hence, we conclude that our findings are robust to potential violations of SUTVA.

Table A6 Efficiency and Resilience Results Excluding Top Five Destinations in Table A5

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Treated (Dallas $\times$ Post), α_3	-34.98523^{***}	-4.89857^{***}	-16.54032^{***}	-0.00886^{***}	0.00231^{***}	-0.00645^{***}
	(0.05649)	(0.05987)	(0.05781)	(0.00016)	(0.00021)	(0.00016)
Treated $\times$ Type II, α_8		-26.90400^{***}			-0.01471^{***}	
		(0.26851)			(0.00071)	
Treated $\times$ Type III, α_9		-70.71403^{***}			-0.02540^{***}	
		(0.10435)			(0.00035)	
Treated $\times$ ACRI, α_{12}			135.37100^{***}			
			(0.21809)			
Treated $\times$ ACRR, α_{12}						-0.11742^{***}
						(0.00186)
Observations	10,896,066	10,896,066	10,844,006	10,892,610	10,892,610	10,840,774
Adjusted R^2	0.78814	0.81446	0.80437	0.01478	0.01645	0.01604

Notes. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Standard errors clustered at the destination level are in parentheses. The dependent variable is the scheduled travel time in Models 1-3 and long delays in Models 4-6. Model specifications are identical to their counterparts in Sections 3 and 4, but the sample excludes passengers traveling from DAL and HOU to ATL, LGA, MCO, TPA, and LAX. We report only the treatment effect coefficients because of space limitations.

B.2. Models By Destination Type

In this section, we estimate separate models for each destination type to check the robustness of the findings we generate from Models 2 and 5. Panel A of Table A8 shows treatment effect

Table A7 Efficiency and Resilience Results Excluding Top 10 Destinations in Table A5

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Treated (Dallas × Post), α_3	-26.52860^{***}	-5.29415^{***}	-15.54226^{***}	-0.00830^{***}	0.00250^{***}	-0.00762^{***}
	(0.06233)	(0.06038)	(0.05940)	(0.00017)	(0.00021)	(0.00017)
Treated × Type II, α_8		-26.70897^{***}			-0.01507^{***}	
		(0.26903)			(0.00071)	
Treated × Type III, α_9		-63.14494^{***}			-0.03130^{***}	
		(0.12738)			(0.00040)	
Treated × ACRI, α_{12}			126.58700^{***}			
			(0.28206)			
Treated × ACRR, α_{12}						-0.13860^{***}
						(0.00212)
Observations	8,880,240	8,880,240	8,828,180	8,877,014	8,877,014	8,825,205
Adjusted R^2	0.80797	0.83114	0.81932	0.01358	0.01596	0.01533

Notes. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Standard errors clustered at the destination level are in parentheses. The dependent variable is the scheduled travel time in Models 1-3 and long delays in Models 4-6. Model specifications are identical to their counterparts in Sections 3 and 4, but the sample excludes passengers traveling from DAL and HOU to ATL, LGA, MCO, TPA, LAX, BWI, BNA, FLL, DCA, and MDW. We report only the treatment effect coefficients because of space limitations.

estimates for Models 2(a), 2(b), and 2(c), which contain approximately 4.4 million, 910,000, and 7.2 million passengers who traveled from DAL and HOU to type I, type II, and type III destinations, respectively. The α_3 coefficient in Model 2(a) shows that the repeal resulted in a 3.15-minute decrease in STT from DAL to type I destinations, on average. In addition, Model 2(b) reveals that the repeal led to a 31-minute decrease in STT from DAL to type II destinations, while Model 2(c) indicates a 78-minute reduction in STT from DAL to type III destinations, on average. Panel B of Table A8 shows our coefficient estimates for Models 5(a), 5(b), and 5(c), which focus on the resilience impact of the repeal on type I, type II, and type III destinations, respectively. In Model 5(a), the α_3 coefficient indicates that the repeal led to a 0.32 percentage point increase in a passenger's likelihood to experience a long delay while traveling from DAL to a type I destination. By contrast, Models 5(b) and 5(c) indicate that the repeal led to 1.33 and 2.0 percentage point decreases in a passenger's likelihood to experience a long delay while traveling from DAL to type II and type III destinations, respectively. Collectively, these estimates are in line with their counterparts in Model 2 of Table 4 and Model 5 of Table 5, showing the robustness of our findings to model specification.

Table A8 Efficiency and Resilience Results by Destination Type

	Panel A: Efficieny Models			Panel B: Resilience Models		
	Model 2(a)	Model 2(b)	Model 2(c)	Model 5(a)	Model 5(b)	Model 5(c)
Treated (Post × Dallas), α_3	-3.14891^{***}	-31.21235^{***}	-78.14379^{***}	0.00322^{***}	-0.01329^{***}	-0.02002^{***}
	(0.06072)	(0.26032)	(0.07183)	(0.00021)	(0.00068)	(0.00024)
Observations	4,381,367	909,555	7,177,862	4,379,220	909,347	7,176,231
Adjusted R^2	0.54613	0.66477	0.62594	0.01475	0.01544	0.01677

Notes. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Standard errors clustered at the destination level are in parentheses. The dependent variable is the scheduled travel time in Models 2(a)-2(c) and long delays in Models 5(a)-5(c). The model specifications of Models 2(a)-2(c) are identical to the model specification of Model 2 in Table 4. The model specifications of Models 5(a)-5(c) are identical to the model specification of Model 2 in Table 5. Models 2(a) and 5(a) focus on type I destinations, Models 2(b) and 5(b) focus on type II destinations, and Models 2(c) and 5(c) focus on type III destinations. We report only the treatment effect coefficients because of space limitations.

B.3. Alternative Time Windows

Our main models focus on the two-year period around the repeal. In this section, we consider two alternative time windows to test the robustness of our findings. The first time window is the 18-month period around the repeal (i.e., January 12, 2014, and July 13, 2015). In this time window, the period between January 12, 2014, and October 12, 2014, represents the pre-repeal period, and the period between October 13, 2014, and July 13, 2015, represents the post-repeal period. The second time window is the one-year period around the repeal (i.e., April 13, 2014–April 12, 2015), where the first six months represent the pre-repeal period, and the last six months represent the post-repeal period. We re-estimate Models 1–6 for each time window. Tables A9 and A10 show the treatment effect coefficients for the one-year and 18-month time windows surrounding the repeal, respectively. These estimates are in line with their counterparts in Tables 4 and 5. Hence, we conclude that our results are robust to alternative time windows.

Table A9 Results for the One-Year Period Surrounding the Repeal

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Treated (Dallas × Post), α_3	-37.19800^{***}	-3.72305^{***}	-13.38315^{***}	-0.00720^{***}	0.00362^{***}	-0.00422^{***}
	(0.07321)	(0.08284)	(0.08385)	(0.00021)	(0.00027)	(0.00021)
Treated × Type II, α_8		-22.49695^{***}			-0.02121^{***}	
		(0.29511)			(0.00077)	
Treated × Type III, α_9		-71.50327^{***}			-0.01894^{***}	
		(0.13234)			(0.00043)	
Treated × ACRI, α_{12}			134.86080^{***}			
			(0.25681)			
Treated × ACRR, α_{12}						-0.09157^{***}
						(0.00265)
Observations	6,098,681	6,098,681	6,085,482	6,096,480	6,096,480	6,083,418
Adjusted R^2	0.78556	0.81738	0.80571	0.01633	0.01698	0.01676

Notes. *p<0.05; **p<0.01; ***p<0.001. Standard errors clustered at the destination level are in parentheses. The dependent variable is the scheduled travel time in Models 1-3 and long delays in Models 4-6. Model specifications are identical to their counterparts in Sections 3 and 4. We report only the treatment effect coefficients because of space limitations.

Table A10 Results for the 18-Month Period Surrounding the Repeal

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Treated (Dallas × Post), α_3	-39.14423^{***}	-4.37059^{***}	-15.77056^{***}	-0.00733^{***}	0.00384^{***}	-0.00437^{***}
	(0.05980)	(0.06832)	(0.06723)	(0.00018)	(0.00025)	(0.00018)
Treated × Type II, α_8		-18.31692^{***}			-0.02088^{***}	
		(0.24415)			(0.00067)	
Treated × Type III, α_9		-74.58518^{***}			-0.02094^{***}	
		(0.10770)			(0.00038)	
Treated × ACRI, α_{12}			129.94050^{***}			
			(0.20393)			
Treated × ACRR, α_{12}						-0.11016^{***}
						(0.00220)
Observations	9,273,022	9,273,022	9,242,420	9,269,816	9,269,816	9,239,355
Adjusted R^2	0.77702	0.80867	0.79596	0.01514	0.01608	0.01590

Notes. *p<0.05; **p<0.01; ***p<0.001. Standard errors clustered at the destination level are in parentheses. The dependent variable is the scheduled travel time in Models 1-3 and long delays in Models 4-6. Model specifications are identical to their counterparts in Sections 3 and 4. We report only the treatment effect coefficients because of space limitations.

B.4. Additional Efficiency Analysis

In this section, we test whether our efficiency results remain valid when we use actual travel times (i.e., scheduled travel times plus travel delays) as a proxy for operational efficiency. To this end, we re-estimate Models 1–3 after replacing the original dependent variable, scheduled travel time (STT), with the new dependent variable, actual travel time (ATT). Table 4 shows the treatment effect estimates for Models 1–3 when the dependent variable is ATT. Specifically, Model 1 reveals that the repeal improves the overall efficiency of Southwest's operations for its DAL-origin passengers. In addition, Model 2 shows that type I (type III) destinations benefit the least (most) from the repeal. Last, Model 3 shows that the destinations with large $ACRI_d$ values benefit more from the repeal. Collectively, these findings imply that using ATT (rather than STT) as a proxy for operational efficiency does not change our main insights into the efficiency impact of the repeal.

Table A11 Efficiency Results Based on Actual Travel Times

	Model 1	Model2	Model 3
Treated (Dallas × Post), α_3	-49.49962^{***}	-3.23685^{***}	-22.10464^{***}
	(0.12600)	(0.13799)	(0.13502)
Treated × Type II, α_8		-44.56756^{***}	
		(0.70563)	
Treated × Type III, α_9		-93.05684^{***}	
		(0.24178)	
Treated × ACRI, α_{12}			150.21630^{***}
			(0.45887)
Observations	12,464,798	12,464,798	12,412,962
Adjusted R^2	0.42116	0.44200	0.43505

Notes. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Standard errors clustered at the destination level are in parentheses. The dependent variable is the actual travel time in all three models. The explanatory variables are identical to their counterparts in Section 3. We report only the treatment effect coefficients because of space limitations.

B.5. Additional Resilience Analysis

In this section, we test the robustness of our resilience results. To this end, we first redefine our binary dependent variable using a two-hour cutoff (rather than the three-hour cutoff we used in our main models). That is, we redefine $Delayed_i$ such that it equals one if passenger i's delay exceeds two hours and zero otherwise. Table A12 shows the treatment effect estimates for Models 4–6 with the new dependent variable. The coefficient estimates of all three models are in line with their counterparts in Table 5.

Next, we return to our original binary dependent variable, which equals one if passenger i's delay exceeds three hours, and change our estimation method. Recall that we use linear regression models in Section 3. Table A13 shows the treatment effect estimates for Models 4–6 when we estimate them using a logit model specification. Although switching from a linear model to a nonlinear one changes the magnitudes of the coefficient estimates, such estimates remain directionally consistent with those in Table 5.

Table A12 Resilience Results Based on Delays Exceeding Two Hours

	Model 4	Model 5	Model 6
Treated (Dallas × Post), α_3	-0.00964^{***}	0.00252^{***}	-0.00733^{***}
	(0.00020)	(0.00029)	(0.00021)
Treated × Type II, α_8		-0.01844^{***}	
		(0.00092)	
Treated × Type III, α_9		-0.02334^{***}	
		(0.00042)	
Treated × ACRR, α_{12}			-0.10343^{***}
			(0.00233)
Observations	12,464,798	12,464,798	12,412,962
Adjusted R^2	0.02074	0.02141	0.02133

Notes. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Standard errors clustered at the destination level are in parentheses. In all three models, the dependent variable is a binary variable that equals one if the delay exceeds two hours and zero otherwise. The explanatory variables are identical to the ones in Section 4. We report only the treatment effect coefficients because of space limitations.

Table A13 Resilience Results Based on Logistic Regression

	Model 4	Model 5	Model 6
Treated (Dallas × Post), α_3	-0.50391^{***}	0.14484^{***}	-0.30603^{***}
	(0.00925)	(0.01911)	(0.01090)
Treated × Type II, α_8		-0.62721^{***}	
		(0.03702)	
Treated × Type III, α_9		-0.90751^{***}	
		(0.02217)	
Treated × ACRR, α_{12}			-3.62528^{***}
			(0.10338)
Observations	12,464,798	12,464,798	12,412,962
McFadden Pseudo R^2	0.08310	0.08567	0.08480

Notes. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Standard errors in parentheses. The only difference between these models and their counterparts in Section 4 is that we estimate these models using a logit specification. We report only the treatment effect coefficients because of space limitations.

We further test the robustness of our resilience results by replacing our binary dependent variable with the actual delay in minutes and estimating quantile regressions. Table A14 shows the treatment effect estimates for Models 4–6 when we estimate the 95th percentile of the delay distribution, which is around three hours in our dataset. The coefficient estimates of all three models are directionally consistent with those in Table 5. Collectively, our findings in this section imply that the resilience impact of the repeal is robust to alternative resilience proxies (e.g., delays exceeding two hours) and estimation methods (e.g., logistic and quantile regression specifications).

B.6. Role of Airline Competition

The airline literature documents that competition on a route affects airline on-time performance (e.g. Prince and Simon 2009, Deshpande and Arıkan 2012, Atkinson et al. 2016). The literature also documents that load factors (the percentage of utilized seats) impact the operational reliability of airline schedules (e.g., Deshpande and Arıkan 2012, Atkinson et al. 2016, Nicolae et al. 2017). Therefore, it is plausible that competition and load factors may affect scheduled travel times and travel delays. To corroborate our main results, we perform additional analyses by controlling for load

Table A14 Resilience Results Based on Quantile Regressions

	Model 4	Model 5	Model 6
Treated (Dallas × Post), α_3	-6.16928^{***}	9.46346^{***}	-8.01504^{***}
	(1.44188)	(1.91725)	(1.76137)
Treated × Type II, α_8		-49.75967^{***}	
		(6.82691)	
Treated × Type III, α_9		-46.95732^{***}	
		(4.14796)	
Treated × ACRR, α_{12}			-115.37560^{***}
			(17.00397)
Observations	12,464,798	12,464,798	12,412,962

Notes. *p<0.05; **p<0.01; ***p<0.001. Standard errors in parentheses. In all three models, the dependent variable is the delay in minutes, while the explanatory variables are identical to the ones in Section 4. We estimate all three models using a quantile regression.

factors and route competition. In this section, we first explain how we calculate our competition and load factor control variables. We then present the regression results of models including competition and load factor as control variables.

We use the Herfindahl-Hirschman Index (HHI) as a measure of competition between airlines in an origin-destination city market. U.S. DOT identifies a city market by a unique identifier (i.e., city.market.ID) in the T-100 Domestic Segment dataset. This unique identifier aims to help researchers consolidate airports serving the same city market. For instance, DAL and DFW are assigned the same city.market.ID, as both airports serve the Dallas-Fort Worth metroplex area. It is natural to expect airlines to compete on origin-destination city markets rather than origin-destination airports. Hence, we compute the market shares of airlines on origin-destination city markets based on the total seat offerings during a month. T-100 dataset enables us to calculate the total monthly number of seats offered between two city markets by the competing airlines, including Southwest. Let c denote an origin-destination city market in our dataset and let $S_{c,a,m}$ denote the total number of seats offered by airline a in market c during month m. Then, the market share of airline a on origin-destination city market c is $MS_{c,a,m} = \frac{S_{c,a,m}}{\sum_{i \in \Psi_{c,m}} S_{c,i,m}}$, where $\Psi_{c,m}$ represents the set of airlines that serve the origin-destination city market c during month m. Accordingly, HHI is calculated as $HH_{c,m} = \sum_{i \in \Psi_{c,m}} MS^2_{c,i,m}$. Similar to our congestion variables, we pick the maximum of the HHI values of the origin-destination airport pairs on a passenger's itinerary as that passenger's HHI value.

Load factor is a frequently used operational performance metric in the airline industry. It measures the percentage of seats filled for an airline on a specific route during a month. The Bureau of Transportation Statistics publishes the U.S. Air Carriers Traffic and Capacity data, which can be used to calculate monthly load factor data for each origin-destination pair. Specifically, we collected data from the T-100 Domestic Segment (U.S. Carriers) dataset, which contains information regarding monthly available seats and the number of enplaned passengers for each origin-destination pair during any month. We computed the monthly load factors by dividing the number of enplaned passengers by the available seats offered on each route. Similar to our congestion variables, for every passenger in our dataset, we choose the maximum of the load factors of the origin-destination pairs on the passenger's itinerary as that passenger's load factor value.

We exclude competition and load factor from our main models because of their mismatch with our unit of analysis (i.e., passenger itineraries). However, we provide the treatment coefficient estimates of models that include monthly competition and load factor as additional control variables in Table A15. In light of the coefficient estimates in Table A15, we conclude that our main inferences remain unchanged when we control for competition and load factors.

Table A15 Models with Airline Competition

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Max. load factor	68.81807^{***}	68.94886^{***}	69.36254^{***}	-0.05206^{***}	-0.05149^{***}	-0.05230^{***}
	(0.21307)	(0.20288)	(0.21032)	(0.00072)	(0.00073)	(0.00073)
Max. HHI	93.66610^{***}	70.14360^{***}	79.04824^{***}	0.03102^{***}	0.02449^{***}	0.02678^{***}
	(0.13900)	(0.14985)	(0.14392)	(0.00036)	(0.00036)	(0.00036)
Treated (Dallas × Post), α_3	-33.28938^{***}	-5.41664^{***}	-13.99740^{***}	-0.00485^{***}	0.00209^{***}	-0.00352^{***}
	(0.05259)	(0.05603)	(0.05607)	(0.00015)	(0.00020)	(0.00016)
Treated × Type II, α_8		-20.26519^{***}			-0.01117^{***}	
		(0.26728)			(0.00071)	
Treated × Type III, α_9		-62.18510^{***}			-0.01574^{***}	
		(0.09640)			(0.00032)	
Treated × ACRI, α_{12}			113.95600^{***}			
			(0.17266)			
Treated × ACRR, α_{12}						-0.07910^{***}
						(0.00183)
Observations	12,468,784	12,468,784	12,416,724	12,464,798	12,464,798	12,412,962
Adjusted R^2	0.78843	0.81002	0.80257	0.01700	0.01770	0.01762

Notes. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$. Standard errors clustered at the destination level are in parentheses. The dependent variable is the scheduled travel time in Models 1-3 and long delays in Models 4-6. The only difference between the model specifications in this table and their counterparts in Sections 3 and 4 is that the models in this table include maximum load factor and maximum HHI as two additional control variables. We report only the treatment effect coefficients along with the coefficients of the new control variables because of space limitations.

B.7. Additional Counterfactual Analysis

In the main counterfactual analysis, we focus on the two-leg itineraries originating from DAL during the post-repeal period. In this section, we expand our sample by including the two-leg itineraries originating from HOU during the post-repeal period. This expansion increases our counterfactual sample size to 1.55 million passengers. We follow the numerical optimization procedure described in Section 5 to reallocate passengers with short itinerary buffers to alternative itineraries. Figure A1(a) shows the changes in the average scheduled travel time and the number of passengers experiencing long delays as a function of γ, where we set $\gamma_r = \gamma$ for all $r \in \mathcal{R}$. In addition, Figure A1(b) illustrates the efficiency-resilience tradeoff by showing the change in the number of passengers experiencing long delays as a function of the change in the average scheduled travel time. The tradeoff patterns depicted in these figures are consistent with their counterparts shown in Figure

5. For example, Figure A1(b) illustrates that increasing the average scheduled travel time by 0.3% leads to an approximately 12% reduction in the number of HOU- and DAL-origin passengers experiencing long travel delays. Hence, we conclude that our numerical optimization procedure is robust to alternative samples.

Figure A1 Efficiency-Resilience Tradeoff for Connecting Passengers Originating from HOU and DAL

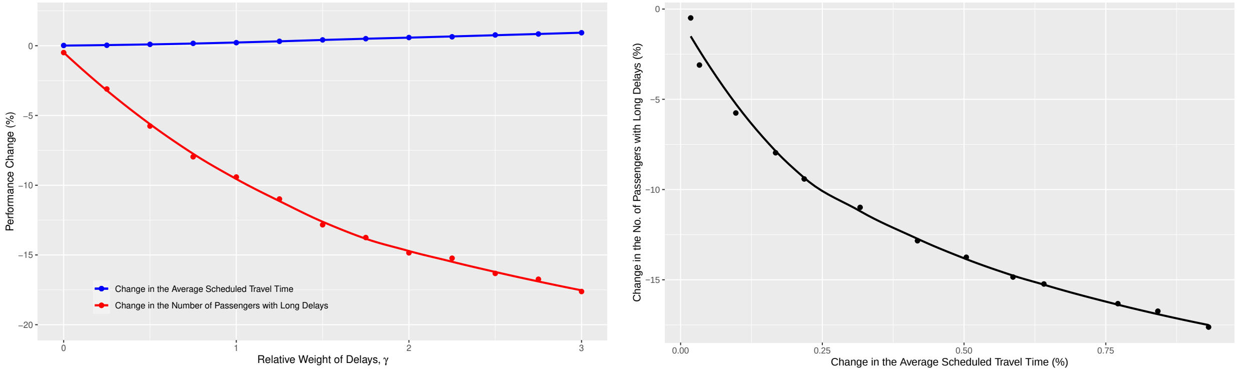

(a) **(b)**

Notes. We set $\gamma_r = \gamma$ for all $r \in \mathcal{R}$ and vary γ between 0 and 3 in increments of 0.25. Figure A1(a) shows the changes in the average scheduled travel time and the number of passengers experiencing long delays as a function of γ. Figure A1(b) uses the same data but places the change in the average scheduled travel time on the horizontal axis and the change in the number of passengers experiencing delays on the vertical axis.

References

Atkinson S, Ramdas K, Williams J (2016) Robust scheduling practices in the US airline industry: Costs, returns, and inefficiencies. *Management Science* 62(11):3372–3391.

Deshpande V, Arıkan M (2012) The impact of airline flight schedules on flight delays. *M&SOM* 14(3):423–440.

Gillingham K, Bollinger B (2021) Social learning and solar photovoltaic adoption. *Management Science* 67(11):7091–7112.

Herring J, Lurkin V, Garrow L, Clarke J, Bierlaire M (2019) Airline customers' connection time preferences in domestic US markets. *Journal of Air Transport Management* 79:101688.

Jo W, Sunder S, Choi J, Trivedi M (2020) Protecting consumers from themselves: Assessing consequences of usage restriction laws on online game usage and spending. *Marketing Science* 39(1):117–133.

Nicolae M, Arıkan M, Deshpande V, Ferguson M (2017) Do bags fly free? An empirical analysis of the operational implications of airline baggage fees. *Management Science* 63(10):3187–3206.

Prince J, Simon D (2009) Multimarket contact and service quality: Evidence from on-time performance in the US airline industry. *Academy of Management Journal* 52(2):336–354.