What's the Difference? Measuring the Effect of Mergers in the Airline Industry

Reed Orchinik *
MIT Sloan

Marc Remer†
Swarthmore College

September 6, 2023

Abstract

We analyze the effect of four US airline mergers using three retrospective techniques: standard difference-in-differences regression, synthetic control, and nearest neighbor matching. We study if these techniques can be reliably applied to study airline mergers, and find a number issues. For example, routes typically specified as controls are potentially impacted by the merger, and the more advanced techniques may be subject to unobservable variable bias. As such, the estimated effect of a merger does not always align in direction or statistical significance across the three methods. We also expand the analysis beyond the average treatment effect on overlap routes. We find that the increase in multi-market contact from each merger caused prices to increase on a large set of non-overlap routes. Lastly, with the aim of providing guidance to antitrust agencies, we analyze the determinants of route-level price effects on overlap routes. We find that even in mergers where the average effect is different from zero (e.g. American/USAir and United/Continental), there is substantial heterogeneity across overlap routes. The route-level effect is predictable based on pre-merger observables, such as the level of HHI.

*MIT Sloan, 100 Main St, Cambridge, MA 02142. Email: orchinik@mit.edu.
†Swarthmore College, 500 College Ave., Swarthmore, PA, 19081. Email: mremer1@swarthmore.edu.

1 Introduction

It is well understood that a merger between competitors can lead to the exercise of market power and a loss of total welfare. However, a horizontal merger may also generate marginal cost efficiencies, which can result in lower prices and increased welfare. Preventing the former mergers and allowing the latter is a central goal of antitrust enforcement. From an empirical standpoint, economists have developed two general sets of tools to help distinguish between pro and anti-competitive mergers. The first set of tools are prospective, such as demand estimation and merger simulation, which are used to forecast the potential competitive impact of a merger. The second set are retrospective, which are used to study the actual impact of a previously consummated merger.

In this article, we implement and analyze three merger retrospective techniques: standard difference-in-differences regression (DID), synthetic control, and nearest neighbor matching. We employ these techniques in order to understand the impact of airline mergers in the United States, and to more generally provide guidance to both antitrust practitioners and researches performing merger retrospectives. We analyze four mergers in the US airline industry: Delta/Northwest, United/Continental, Southwest/AirTran, and American Airlines/USAir. We first focus on the effect of the mergers on non-stop, overlap routes. We find that airline mergers have a heterogeneous effect on prices. However, estimation results can be sensitive to the choice of econometric method. We study potential sources of the disparate results across the methods, and find that each may be prone to bias when applied to the airline industry.

The analysis is then broadened to include the impact of the merger on routes where only one merging firm competed prior to the merger. We find that these routes experienced an increase in price, likely to due to the increased incentive to coordinate with non-merging firms. This finding demonstrates that analyzing price changes only in overlap markets is inappropriate in industries prone to coordination. Furthermore, including these routes in the control group leads to biased effect estimates. More generally, it shows that industries prone to coordination may be a poor fit for implementing DID techniques to perform merger retrospectives.

Despite the issues we identify with each method, on routes in which both airlines competed head-to-head prior to the merger, we find that the American/USAir merger is likely to have generated lower prices. The United/Continental merger resulted in higher prices on overlap routes. The impacts of the Delta/Northwest and Southwest/AirTran mergers on overlap routes are much less clear. Indeed, we estimate that the Southwest/AirTran merger significantly increased prices (8% to 9%) using DID regression, but significantly decreased prices (-9% to -19% in the first post-merger year) using nearest neighbor matching. A common finding across all mergers is that the estimated price effects are sensitive to either the applied methodology or reasonable choices made within a given methodology.

We identify multiple reasons why standard OLS DID estimation of merger price effects on overlap route may be a poor fit for the US airline industry. For example, in each merger there are a very small number of overlap routes relative to control routes, overlap routes are structurally different than controls, and the parallel trend test often does not hold. Each of these factors suggest that synthetic control or nearest neighbor are more likely to produce unbiased estimates of the average treatment effect.

We show that, according to standard metrics, synthetic control and nearest neighbor create more suitable control groups than OLS DID when performing merger retrospective analysis in the airline industry. However, we demonstrate that, in practice, these methods are also sensitive to choices made by the researcher, and that OLS DID, synthetic control, and nearest neighbor matching should not be expected to produce similar results. Furthermore, in the context of merger retrospectives, we derive an easily implementable metric to test whether synthetic control and nearest neighbor matching properly control for unobservable determinants of pre-treatment outcomes. We are unable to rule out that synthetic control and nearest neighbor are subject to unobservable variable bias.

In order to draw sharper conclusions about the impact of airline mergers, we undertake two additional sets of analyses. As has been developed in previous research, the US airline industry is prone to coordination through the mechanism of multi-market contact.[1] A merger in the airline industry creates a sharp increase in multi-market contact on routes where only a single merging firm competes pre-merger. In each merger, there are approximately 10 times more of these routes than overlaps. As such, to the extent that coordination increased on these routes post-merger, it could lead to significantly more welfare loss than price increases on overlap routes. Our results show that in each of the four mergers, prices increased on these routes due to the merger-induced increase in multi-market contact.

Our final set of results are aimed at providing guidance to antitrust authorities in predicting route-level price effects on overlap routes. They also help to explain the heterogeneous impact of mergers on these routes. To determine which pre-merger observables are best able to explain route-level price changes, we implement a cross-validated lasso regression and find that pre-merger measures of HHI and a small set of additional route characteristics are able to explain post-merger outcomes on overlap routes. In particular, higher pre-merger HHI

[1] See, for example, Ciliberto and Williams (2014) and Ciliberto, Watkins, and Williams (2019).

leads to higher post-merger prices, and the substitutability of connecting flights is inversely related to price changes. The findings can largely, but not entirely, be explained by a model where markets are prone to price coordination.

In April, 2019, the Federal Trade Commission held public hearings discussing the importance of merger retrospectives.[2] In the hearings, the extent to which merger retrospectives can be used to understand the outcomes of previous mergers and generalized to prospective mergers was debated. While the exact value of merger retrospectives is an open question, there is no doubt that accurately measuring the impact of horizontal mergers will greatly benefit antitrust policy. Indeed, since these hearings, the Federeal Trade Commission has implemented a "Merger Retrospective Program," to assess the impact of consummated mergers.[3] By implementing and assessing three leading merger retrospective techniques on real-world consummated mergers, we demonstrate both the potential value and limitations of ex-post evaluation of mergers.

Understanding the impact of consolidation in the US airline industry is particularly important. In the past 15 years, horizontal mergers have left the industry with only three legacy airlines – American, Delta, and United – and one predominant low-cost carrier – Southwest. In its lawsuit to block the American Airlines/USAir merger, the US Department of Justice noted that this merger wave has resulted in higher fares and reduced service between city pairs.[4] More recently, the industry has been under investigation by the DOJ for potentially coordinating capacity discipline in an effort to increase prices,[5] and the DOJ is currently suing to block the merger of two low-cost carriers – JetBlue and Spirit.[6] Given the recent consolidation and airlines' importance to the overall economy, it is imperative to understand precisely if mergers in the industry have resulted in price increases and lower consumer welfare.

While accurate retrospective measurements are important for building general knowledge about mergers, they are also used to gauge the accuracy of merger prediction tools. A number of articles first estimate a structural demand model using pre-merger data and then simulate a merger. The accuracy of the simulations are then measured relative to the DID estimates.[7]

[2]`https://www.ftc.gov/news-events/events-calendar/ftc-hearing-14-merger-retrospectives`

[3]For program details, see `https://www.ftc.gov/policy/studies/merger-retrospectives`.

[4]See paragraph 35 of DOJ's complaint against American Airlines and USAir, which is available at `https://www.justice.gov/atr/case-document/file/514531/download`.

[5]See Ciliberto, Watkins, and Williams (2019) and cites therein.

[6]For the DOJ's complaint, see `https://www.justice.gov/d9/press-releases/attachments/2023/03/07/001_-_complaint_jb_0.pdf`.

[7]See, for example, Peters (2006), Weinberg and Hosken (2013). Garmon (2017) uses the synthetic control method to form a baseline estimate of merger price effects.

This analysis assumes that the merger retrospective accurately measures the true effect of the merger. We demonstrate that this may currently place too high of a burden on commonly used retrospective techniques to accurately estimate merger effects. The three retrospective methodologies analyzed in this article may produce estimates of greatly different magnitude, and even sign. As best practice, we encourage researchers to implement either nearest neighbor or synthetic control alongside standard DID when conducting a merger retrospective analysis.

The remaining paper is organized as follows. Section 2 reviews the related literature, Section 3 details the three estimation methodologies. Section 4 summarizes the data, details pre-merger pricing trends, and how well different methodologies fit the pre-merger data. Section 5 presents the estimation results for overlap routes, Section 6 present the multi-market contact analysis, Section 7 analyzes route-level effects, and Appendix Section B develops the test for unobserved variable bias in synthetic control and nearest neighbor matching.

2 Literature Review

There is a rich and growing literature that uses DID regressions to estimate the competitive impact of consummated mergers. For example, the price effect of horizontal mergers between health insurance providers (Dafny, Duggan, and Ramanarayanan, 2012), beer companies (Ashenfelter, Hosken, and Weinberg, 2015), and academic journals (McCabe, 2002) have all been estimated using DID. In a closely related literature, the DID estimated effects of a consummated merger are compared to the predictions of a structural analysis estimated on pre-merger data. Using this approach, the airlines industry (Peters, 2006), hospital markets (Garmon, 2017), motor oil and maple syrup products (Weinberg and Hosken, 2013), and retail gasoline markets (Houde, 2012), among others, have been analyzed. Ashenfelter, Hosken, and Weinberg (2014) perform a comprehensive review of the merger retrospective literature, and there are two notable findings. Almost all retrospectives use standard DID regression, and mergers in oligopoly markets tend to generate higher prices. Two studies that use more advanced DID estimators (such as synthetic control and propensity score mathcing) in merger retrospectives are Hosken, Olson, and Smith (2018) and Allain, Chambolle, Turolla, and Villas-Boas (2017). We differ from these studies by focusing on the extent to which the different methodologies produce different results, their sensitivity to different specifications, and developing a test for unobserved bias. We also contribute to this literature by estimating the effects of four separate mergers in the airline industry, and comparing the estimates across several different methodologies.

There is a long history in economics of analyzing the US airline industry, in general, and horizontal mergers within the industry. Following the Airline Deregulation Act of 1978, there was a large wave of industry consolidation. Borenstein (1990) evaluated the 1986 mergers of TWA/Ozark and Northwest/Republic by comparing prices on overlap routes to those on unaffected and similar length routes, and finds that prices were unchanged in the former merger and increased by almost 10% in the latter. Using a set of control routes distinct from Borenstein (1990), Morrison (1996) analyzes the same two mergers (in addition to USAir/Piedmont) and finds that over a similar time frame used in Borenstein (1990) TWA/Ozark prices increased and Northwest/Republic decreased. Kim and Singal (1993) use standard DID regressions to estimate the price effect of 14 airline mergers that occurred from 1985 to 1988, and find that market power effects outweighed efficiency gains. More recently, Huschelrath and Muller (2013) analyzes six US airline mergers from the late 1990's up through the Delta/Northwest merger in 2008 and finds that prices increased by 3%-6%, depending upon the post-merger time frame considered. Das (2019) uses standard DID and finds that the American/USAir merger led to a decrease in prices and no change in flight frequency. Lazarev, Nevo, and Town (2021) analyze the effect of United/Continental and American/USAir, focusing on the extent to which welfare consequences of the mergers can be identified from DID estimates of price and quantity effects, and find results directionally similar to ours.

In terms of mergers analyzed, Carlton, Israel, MacSwain, and Orlov (2019) is most closely related to this article. It uses standard DID to estimate the effects of the three most recent US legacy airline mergers: Delta/Northwest, United/Continental, and American/USAir. Following the advice of Mehta and Miller (2012), Carlton, Israel, MacSwain, and Orlov (2019) take seriously the potential sensitivity of estimates to the assumed control group. The analysis includes a number of robustness checks wherein the merger effects are measured relative to different sets of control routes pre-selected by the authors. The study finds that Delta/Northwest and American/USAir led to significant price decreases on overlap routes, and there was no signficant change in prices due to the United/Continental merger.

The empirical analysis in this article differs from Carlton, Israel, MacSwain, and Orlov (2019) along a number of important dimensions. First, we implement two additional empirical methodologies that allow the data to guide the selection of control routes, rather than relying exclusively on our own judgment. While synthetic control or matching estimators are frequently utilized in applied microeconomics, they are seldom employed in industrial organization.[8] Second, a primary aim of this article is to compare the DID estimates to the newer techniques that weight control units to more closely match treated units. Our results demonstrate that synthetic control and nearest neighbor estimators may lead to very different findings than standard DID regressions. We additionally study a merger between two low cost carriers (Southwest/AirTran), as there may be a market impact unique from legacy mergers. Finally, we also analyze each merger at the route-level and find important differences across routes, within the same merger.

Finally, our paper relates the literature on coordination via multi-market contact. The theory was first developed in Bernheim and Whinston (1990). In two separate studies of the airline industry, Ciliberto and Williams (2014) and Ciliberto, Watkins, and Williams (2019) provide evidence that multi-market contact leads to higher prices and that pricing patterns are generally consistent with collusive behavior. We contribute to this literature by exploiting a plausibly exogenous increase in multi-market to study its causal impact on airline prices. Additionally, we connect this finding to the merger retrospective literature to show how such coordination can lead to biased estimates on routes where merging parties compete head-to-head prior to the merger.

3 Methodology and Merger Retrospective Overview

To analyze the effect of airlines mergers on prices, we implement three separate methodologies: difference-in-differences OLS regression (DID), synthetic control, and nearest neighbor matching. The aim of each technique is to infer the effect of the merger by comparing prices on overlap routes relative to a set of control routes unaffected by the merger. The primary distinction between the methods is in how they select and weight control routes. In this section, we give a high level overview and comparison of each methodology in the context of merger retrospectives. In Appendix section A we provide more technical detail.

A merger retrospective analysis using difference-in-differences-style methods requires panel data. Additionally, the analysis typically leverages two industry facts: (i) firms compete in many distinct geographic markets and (ii) there is variation in market structure across geographic markets. The "treatment" is then defined as the exogenous change in market structure due to the merger, which occurs only in markets where both firms competed pre-merger. The treatment effect is the causal impact of the merger on prices (or another outcome variable) in the treated markets.

[8] A notable exception is Deryugina, MacKay, and Reif (2020), which uses nearest neighbor matching to estimate the effect of electricity regulations on market demand.

3.1 Difference-in-differences OLS Regression

DID regression is the most common technique used in merger retrospective analysis. To identify the causal effect of the merger, OLS regression is employed along with assumptions on the evolution of outcomes in "treated" and "control" markets. Specifically, an equation of the following form is estimated:

$$y_{mt} = \alpha + \beta D_{mt} + \gamma_m + \lambda_t + \epsilon_{mt}. \tag{1}$$

Equation (1) is specified at the market-level, m, and is estimated using panel data that varies over time, t.[9] Here, y_{mt} is an outcome of interest, γ_m is a market fixed effect, λ_t is a time fixed effect, and ϵ_{mt} is an iid normally distributed error term. Assuming the post-merger period begins at time r and M is the set of affected markets, then D_{mt} equals one for all $t \geq r$ and $m \in M$, and zero otherwise. The coefficient β therefore measures how the outcome, y_{mt}, changes post-merger in affected markets *relative* to unaffected markets.

In the airline industry, a market is often defined as a route between specific origin and destination airports.[10] We restrict our attention to non-stop routes, as this market definition has been employed by the US Department of Justice in its antitrust analysis and makes our results comparable to other US airline merger retrospectives (e.g. Werden, Joskow, and Johnson (1991) and Carlton, Israel, MacSwain, and Orlov (2019)). Routes affected by the merger are those where the merging airlines both operate regular non-stop service prior to the merger.[11]

In addition to the standard OLS assumptions, there are more conditions that must hold for β in equation (1) to be an unbiased estimate of the average treatment effect of the merger on routes M. For the purposes of this article, we focus on assumptions specific to merger retrospectives and refer readers to Wooldridge (2002) for a full treatment of DID estimation.[12] First, within each market, it must be that the merger is an exogenous event.

[9]Equation (1) can also be estimated using firm-level data. However, due to institutional details described below and ease of comparing across methods, we aggregate the data to the market level.

[10]See, for example, Ciliberto and Tamer (2009) and Berry and Jia (2010). Alternatively, some studies define a market using city-pairs rather than airport-pairs. Comparing these market definitions is beyond the scope of this paper.

[11]Dix and Orzach (2021) find that airline mergers may result in price increases on connecting routes as well. Thus, our estimates may be a lower bound on the potential price impact of the airline mergers.

[12]A recent literature on two-way fixed effects demonstrates that bias may arise in DID if treatment is staggered over time (Baker, Larcker, and Wang (2022)). These findings do not apply in our setting, as treatment is applied simultaneously across routes. Sloczynski (2022) finds that the standard weighting scheme for OLS DID may be lead to biased estimates in many instances. However, when the average treatment effect on the treated (ATT) is the object of interest, and the number of treated units is small relative to the number of control units, then standard weighting is appropriate. This is precisely the case in airline merger

This assumption is violated if, for example, two airlines merge because they expect future demand to increase on their overlap routes relative to non-overlap routes.

Second, it is critical that unaffected routes provide a baseline for what would have occurred absent the merger. Specifically, there must be no systematic or time-varying differences between the treatment and control groups that are not accounted for. A potential concern in the airline industry is if the merger leads to an increase in system-wide anti-competitive effects, perhaps due to an increased ability to sustain price coordination.[13] If this were to occur then non-overlap routes would provide a poor baseline against which to measure the impact of the merger, as prices on these control routes would also increase due to the merger. In results presented below, we demonstrate that this likely occurs and, thus, impacts estimated effects in the airline industry.

If coordination does not arise in non-overlap markets due to a merger, then standard methods for addressing cross-sectional or time-varying differences may be applied. In equation (1), market fixed effects control for time-invariant differences across routes, and time fixed-effects control for time varying factors common across all routes. While this fixed-effect structure helps account for many unobservables, they are not a panacea. If, for example, population or income is growing faster on treatment routes then the estimate of β may be biased upwards. More specifically, DID assumes that treatment and control groups follow a parallel trend prior to the treatment, and would have continued on parallel trends but for the treatment. In other words, the treated and control groups must have the same data-generating process conditional on regression controls. In a setting with multiple pre-treatment periods, such as in our analysis, this assumption may be tested directly.

Even if the parallel trend assumption holds, further refinements are often made to the control group in DID analysis. In the context of horizontal mergers, the control group may be limited to markets with a similar competitive environment as the pre-merger overlap routes. In airline mergers, if the treatment group only includes routes with large hubs at either end, then it may be appropriate to exclude routes between small regional airports. These decisions, however, are typically guided by economic intuition rather than a precise economic model.

Synthetic control and nearest neighbor matching offer a data-driven approach to selecting a control group that is potentially a better reflection of the but-for world of the treatment group. These techniques, at least in principle, decrease the amount of discretion or economic

retrospectives.

[13]The network structure of the airline industry and multi-market contact are two potential reasons why the airline industry may be prone to collusion.

intuition needed to select an appropriate control group.

3.2 Synthetic control estimation

Synthetic control was first developed to assess policy effects in a setting with relatively few treatment and control units (Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainmueller (2010)). In such an environment, it may be that each control unit serves as a poor proxy for the treated unit(s). However, a weighted average of the control units may produce a "synthetic" control that provides a close match to the treated unit prior to the policy enactment. The treated unit is then compared to the synthetic control during the post-treatment period to measure the treatment effect.

In merger retrospectives, there are typically multiple treated units and a large number of potential control units. The synthetic control method offers the potential to, a priori, remain agnostic about the relevance of each control observation. By selecting a set of characteristics, such as lagged prices or the number of pre-merger competitors, a synthetic control can be constructed by solving for a set of weights that minimize the distance between the pre-treatment characteristics of the treated unit and a weighted average of the control units.

In this article, we apply the extension of Abadie, Diamond, and Hainmueller (2010) developed in Cavallo, Galiani, Noy, and Pantano (2013) for multiple treated units. Here, we give a brief overview of the methodology and discuss issues pertinent to merger retrospectives. More technical details are provided in Appendix section A.1. Let the pre-merger time periods be time $t = 1, .., r-1$ and the post-merger period to be $t = r, ..., T$. Then, define Y_{mt}^N as the price that would be observed on route m at time t if no merger occurred, and let Y_{mt}^I be the price if the merger did occur. The effect in market m of the merger at time $t \geq r$ is then defined as $\alpha_{mt} = Y_{mt}^I$ - Y_{mt}^N. In markets affected by the merger, however, Y_{mt}^N is unobserved for $t \geq r$, and therefore must be estimated.

To estimate Y_{mt}^N, a set of control markets, C, and a set of matching characteristics, Z, which typically include pre-treatment outcome variables, must be specified. Then, if a set of weights, w_c^*, exist such that applying the weights to the selected characteristics on the control routes yield values close to those on a the treated route, then the following is an unbiased estimate of the treatment effect:

$$\widehat{\alpha}_{mt} = Y_{mt}^I - \sum_{1=c}^{C} w_c^* Y_{ct}. \tag{2}$$

To implement the synthetic control method, we solve for a synthetic control unit for each route where the merging firms both offer non-stop service prior to the merger. Then, to calculate the average treatment effect of a merger, in each post-merger time period, we take the average of $\widehat{\alpha}_{mt}$ across treated routes: $\bar{\alpha}_t^a = \sum_m^M \widehat{\alpha}_{mt}$. To test for significance, we follow the permutation tests developed in Cavallo, Galiani, Noy, and Pantano (2013), the details of which are provided in the Appendix A.1.

While the synthetic control method allows the researcher to remain agnostic about the relevance of each control unit, there is still some discretion afforded. The variables upon which to base the synthetic control, Z, need to be specified. This choice is analogous to choosing dependent variables in a DID regression. Kaul, Klößner, Pfeifer, and Schieler (2015) demonstrate that the common practice of using the full set of pre-intervention outcomes in conjunction with additional covariates to construct the synthetic control is likely to be inappropriate, as the the additional covariates typically receive no weight despite their relevance. We therefore estimate two synthetic control specifications. The first only matches on pre-treatment outcome variables (i.e. prices), as is common in the literature. We estimate a second that includes additional covariates and a subset of pre-treatment prices. To select the optimal set of pre-treatment prices we implement a cross-validation routine detailed in Appendix A.1.1.

Moreover, to the best of our knowledge, there is no definitive metric to determine whether, prior to treatment, the synthetic control unit is sufficiently close to the treated unit. Abadie, Diamond, and Hainmueller (2010) exclude control units where the mean square error is more than 20 times the treated unit; Cavallo, Galiani, Noy, and Pantano (2013) informally gauges the goodness of fit of the synthetic control relative to the average value of the dependent variable. Average treatment effect estimates may therefore be sensitive to the synthetic control "convergence" tolerance, especially when there are a small number of treated units.

One clear advantage of synthetic control, however, is that the parallel trend assumption will be satisfied at least as well as in DID. Still, as mentioned above, if control routes are affected by the merger, perhaps due to increased price coordination, then synthetic control will also produce a biased estimate of the merger effect. Furthermore, under certain conditions, synthetic control (and nearest neighbor), may be subject to unobservable variable bias even if the parallel trend test passes. We develop a metric in context of merger retrospectives that helps understand the extent to which unobservable bias may be present, even if the synthetic control and treated units perfectly match in the pre-merger time periods.

3.3 Nearest neighbor matching

Nearest neighbor matching is very similar to the synthetic control method. We apply the methodology developed in Abadie and Imbens (2006, 2011), the technical details of which are provided in Appendix section A.2. Again, a set of control markets, C, and a set of matching characteristics, Z, must be selected by the researcher. Also, let Y_{mt}^N be the price on route m at time t if no merger had occurred, and let Y_{mt}^I be the observed price on a route where the merger did occur. To estimate Y_{mt}^N on treated routes, a finite set of "nearest neighbor" control routes is selected such that applying equal weight to each neighbor produces values of Z close to those on the treated route. The researcher must pre-specifiy how many nearest neighbors, $\mathcal{J}_N$, will be used to estimate Y_{mt}^N.

The estimated treatment effect on route m is $Y_{mt}^I - \widehat{Y}_{mt}^N$, where:

$$\widehat{Y}_{mt}^N = \frac{1}{N} \sum_{j \in \mathcal{J}_N(m)} \{Y_{jt} + \widehat{\mu}_t(x_m) - \widehat{\mu}_t(x_j)\}. \tag{3}$$

In equation (3), $\widehat{\mu}_t$ is a bias correction term developed in Abadie and Imbens (2011) to account for nearest neighbor estimates being inconsistent when matching occurs with more than one continuous variable. The bias correction term, $\widehat{\mu}_t$, is estimated by regressing the outcome variable on the set of matching covariates, and then using the predicted value for each route. The average treatment effect at time t is then:

$$\widehat{\tau}_t = \frac{1}{M} \sum_{m=1}^{M} Y_{mt}^I - \widehat{Y}_{mt}^N, \tag{4}$$

where M is the number of routes that are affected by the merger. Standard errors are derived as in Abadie and Imbens (2011).[14]

Nearest neighbor matching is similar to the synthetic control method in that it allows the data to determine the relevance of each control unit while still leaving important choices to the researcher. As with synthetic control, the set of matching characteristics must be selected. Abadie and L'Hour (2021) demonstrate that both synthetic control and nearest neighbor are special cases of the same class of estimator. As such, for sake of comparison, we estimate nearest neighbor using the same set of covariates: lagged outcome variables and the route characteristics we detail below.[15] Additionally, the number of neighbors must be

[14] See those respective articles for more detail.

[15] Lagged outcomes are commonly used as covariates in nearest neighbor matching. See, for example, List, Millimet, Fredriksson, and McHone (2003), Imbens (2015), and Deryugina, MacKay, and Reif (2020). Given

chosen.

A potential disadvantage to nearest neighbor matching is that a lot of data from the control units are discarded. Where synthetic control may potentially place positive weight on each control unit, nearest neighbor constrains all but N units to have zero weight for each treated unit. As matching is typically done with replacement across treated units, this may result in a majority of data being ignored in estimation. Of course, this may be viewed as an advantage, as all but the most relevant control units are considered. Also, in practice, synthetic control typically results in a sparse weighting matrix where only a small number of control units receive positive weight. Finally, as with synthetic control and DID, if control routes are systematically impacted by the merger then nearest neighbor will yield biased estimates of the treatment effect.

4 Data and Summary Statistics

4.1 Data overview

The data used in this article comes from the US Department of Transportation's Origin and Destination Survey database (DB1B), and is commonly used in the economic analysis of the airline industry.[16] The data is a 10% sample of all domestic airline tickets, and include the quarter of travel, origin and destination airport, airline, number of stopovers, and the price paid for each ticket in the sample.

To conduct the analysis, we restrict the data to only include routes with regular traffic[17] and, similar to previous studies[18] we remove tickets with very low or high prices. We specify routes to be non-directional for a number of reasons. First, recent research demonstrates that airlines almost always set the same price for the same seat, whether or not the round trip is A-B-A or B-A-B . Second, a given airline almost always faces the same set of competitors, regardless of whether the flight is A-B or B-A. Third, the US government often defines non-directional routes when performing its analysis.[19] Thus, specifying directional routes

the results in Abadie and L'Hour (2021), the conditions for using lagged outcomes as covariates in nearest neighbor matching are the same as in synthetic control, where it is standard practice.

[16]See, for example, Goolsbee and Syverson (2008) and Berry and Jia (2010).

[17]Following Carlton, Israel, MacSwain, and Orlov (2019), we only consider routes with an average of 20 passengers per day. Because the data is a 10% sample, this is a cutoff of approximately 200 passengers per day.

[18]Round trip equivalents below below \$50 and above \$2,000 are removed from the sample. Ciliberto, Watkins, and Williams (2019) similarly restrict their data.

[19]See, for example, United States Government Accountability Office's analysis of airline competition. `https://www.gao.gov/assets/670/664060.pdf`.

avoids creating near-duplicate observations and is consistent with how the US government has defined airline markets.

In accordance with antitrust practice and much of the academic literature, we restrict the data to only include non-stop routes. As in Carlton, Israel, MacSwain, and Orlov (2019), we aggregate the data up to the route-quarter level and take the quantity-weighted average price as the primary outcome variable. The analysis is performed at the route-level rather than the airline-route level for two reasons. First, it is more natural to construct synthetic controls and perform matching at the route level, rather than matching on airline-route observations. Second, in response to an exercise of market power, non-merging firms will typically change their prices. Measuring price at the route level accounts for these reactions and gives a more complete accounting of the change in prices paid by consumers.

The data allow us to identify the number of passengers buying tickets at a given price, on each airline on each route. This information is employed to construct variables on which to base the synthetic control units and nearest neighbor matches. Specifically, we create quantity-based HHI's, which capture the level of competition on a route. We also use the total number and quantity share of legacy carriers[20] as an additional competitive metric upon which to match routes. Previous research, such as Brueckner, Lee, and Singer (2013), demonstrates that legacy carriers create less competitive pressure than low-cost carriers (LCC's), and therefore it may be important to separate out legacy carrier impact when matching routes.

4.2 Selecting Treatment and Control Routes

We focus on four US airline mergers: Delta/Northwest, United/Continental, Southwest/AirTran, and American/USAir. The Southwest/AirTrain merger combined low-cost carriers and the other three mergers were between legacy airlines. For each merger, the analysis includes the eight quarters before and after the merger, and we drop from the sample the quarter in which the merger was consummated. We focus on routes that are most likely to be impacted by the merger. As such, we define a route as an "overlap" if both merging airlines had at least a 10% market share in all eight quarters prior to the merger being consummated. Relative to the total number of routes each airline flies, the merging airlines overlap with non-stop service on a small number of routes.

Selecting the control routes against which to measure the impact of an airline merger is

[20]The legacy carriers in the sample are, American Airlines, Continental, Delta, Northwest, United, and USAir.

critical to the analysis. Before implementing any of the methodologies, similar to Carlton, Israel, MacSwain, and Orlov (2019), we discard routes with passenger counts that are too dissimilar from the overlap routes. For each merger, we pool the set of overlap routes and calculate the minimum and maximum number of observed passengers in any quarter. We then drop any route such that, in any quarter, the total number of passengers is 10% less than the minimum or 10% greater than the maximum observed on the overlap routes. In theory, synthetic control and nearest neighbor matching should not place weight on these dropped routes if they are too dissimilar from the treated routes. Still, we apply these filters prior to implementing all three methods so that results are comparable across methodologies and to previous DID studies.

Finally, we split the control routes into two categories: (i) neither merging firm offered non-stop service in the two years prior to the merger and (ii) exactly one merging firm had at least a 10% share in each of the 8 quarters prior to the merger. Any non-overlap route that does not fall into these two categories is dropped from the analysis. We refer to control routes with no merging firm pre-merger as control group 0, and control routes with one merging firm pre-merger as control group 1.

We create two control groups in order to account for horizontal mergers having two potential countervailing effects on price: increased market power and marginal cost efficiencies. Measuring the merger effect relative to routes where neither firm has a pre-merger presence facilitates an estimate of the net impact of the merger.[21] That is, overlap routes potentially include both market power and efficiency effects. If routes where one merging carrier operated pre-merger were included in control group 0, then potential cost savings would be included in both the treatment and control, which would lead to a biased estimate of the net merger effect.

Estimating the merger effect relative to routes where one carrier was present pre-merger allows us to separate out the market power effect from the cost efficiency effect. If cost savings are realized equally across routes and pass-through is, on average, the same across control and treated routes then this isolates the market power effect from the cost saving effect. For example, suppose we find a negative impact on overlap routes relative to control group 0, but a positive impact relative to control group 1. This would suggest that cost efficiencies outweighed the exercise of market power (negative relative to group 0), but that there was still a market power effect (prices increased relative to group 1). In other words, the pass-through of marginal cost efficiencies was lower on overlap routes than on routes with

[21]This assumes that routes where neither firm had a pre-merger presence were unaffected by the merger. This assumption would be violated if, for example, the merger facilitated network-wide price coordination.

only one merging firm, due to greater post-merger market power on overlaps.

This reasoning relies on two important assumptions. First, it assumes that the merging firm has, on average, the same cost pass-through rate on overlap and control group 1 routes. This may not be the case if, for example, the curvature of demand is systematically different on the two sets of routes. Second, it assumes that overlap and control group 1 routes experience the same marginal cost efficiencies. This would be violated if overlap routes experience greater opportunities for cost savings. Given these assumptions, however, control group 1 has the potential to isolate the market power effect from the cost savings effect. More generally, it is prudent to remove the merging parties from control group 0, so that a clean test of the net effect of the merger can be obtained.

4.3 Summary Statistics and Pricing Patterns

For each merger, Table 1 summarizes the treatment and control routes and the competitive environment before and after the merger. The number of overlap routes range from 5 (Delta/Northwest) to 13 (Southwest/AirTran). While this is a small number of routes compared to both control groups, the change in market structure on overlap routes is dramatic. The average change in HHI on overlap routes is more than 3,000 for all four mergers;[22] the 2010 US Horizontal Merger Guidelines state that mergers resulting in a change in HHI greater than 200 are "presumed to be likely to enhance market power." Moreover, the affected routes are "highly concentrated markets" according to the Guidelines, with an average pre-merger HHI of more than 3,000 for each merger.[23] The small number of overlaps and high changes in market concentration are a result of the hubbing system, whereby airlines cluster their flights at particular airports. One potential difficulty with merger retrospectives is that only mergers that are not successfully challenged by antitrust agencies are able to be studied, which may result in only mergers with, at most, small price effects. Here, however, the overlap routes appear prone to the exercise of market power. Yet, since they are a small percentage of overall routes, the antitrust agencies may choose not to challenge the entire merger based on these markets, especially if network-wide efficiencies are expected.

Table 1 also highlights that overlap routes are, on average, structurally different from the control routes. Prior to the merger, overlap routes tend to be less concentrated than the typical route. This is because, by definition, overlap routes have at least two airlines with at least 10% market share, and most other routes tend to be dominated by a single airline.

[22]We use the average quarterly passenger share on each route in the year prior to the merger to calculate the change in HHI.

[23]For access to the Guidelines see, `https://www.justice.gov/atr/file/810276/download`.

Furthermore, overlap routes in legacy (low-cost carrier) mergers tend to have a much higher (lower) overall legacy share compared to control routes. In total, the pre-merger market structure of overlap and control routes are systematically different.

Indeed, the control routes used in the analysis have an average pre-merger HHI ranging from about 6500 to 8500. As such, these direct routes tend to either have one or two firms providing regular service. For the subset of control routes that are not monopolies, price coordination may be of concern. Ciliberto and Williams (2014) and Ciliberto, Watkins, and Williams (2019) find empirical evidence of tacit collusion in the airline industry that is facilitated by multi-market contact. We demonstrate below that each merger led to a substantial increase in multi-market contact on control group 1 routes. We therefore investigate and find evidence that this increased incentive to collude led to higher post-merger prices. These routes therefore likely serve as a poor baseline against which to measure price effects on overlap routes. Control group 0 routes, however, do not experience an increase in multi-market contact due to the mergers and provide a better baseline against which to measure the mergers' effect on prices.

Table 2 summarizes prices and passenger counts before and after each merger, separately by treatment and control groups. It also presents t-tests of differences between pre and post-merger values for each variable. Despite dropping control routes with passenger counts more than 10% outside the range of overlap routes, there still exist substantial differences between overlap and control routes. On average, passenger counts in control group 0 are 30% to 40% less than on overlaps. There is not a clear pattern, however, in how price levels differ between treatment and control groups in the pre-merger period. For example, in United/Continental and American/USAir prices are higher on overlap routes, in Southwest/AirTran prices are much lower in the overlap group, and in Delta/Northwest overlap average prices are between the two control group prices. In sum, on average, it does not appear that treatment and control groups are observationally equivalent, suggesting that synthetic control or nearest neighbor may be more appropriate than DID.

For three of the four mergers, prices on overlap and control routes move in the same direction following the merger, indicating that larger industry trends dominate any price effects specific to overlap routes. This highlights the importance (and potential difficulties) of carefully implementing DID methodologies and, in particular, selecting a control group that provides a benchmark against which to measure the impact of the merger.

Further highlighting this fact is that we don't see any major price movement on overlap routes relative to control groups following any of the mergers. Figure 1 depicts quantity-weighted average prices on treatment and control routes for each merger. Figure 9 presents an index of prices on overlap (and control group 1) routes normalized by prices on control group 0 routes. Both figures demonstrate there is no major, sustained price change following the mergers. In Southwest/AirTran, prices on overlaps jump in the second quarter following the merger, but return to baseline two quarters later. The pricing pattern on these overlaps appears to be dominated by seasonal movements rather than merger-related changes. In American/USAir, prices on overlaps fall relative to control group 0 routes, but they appear to be falling prior to the merger, and the trend continues after the merger is consummated. While the "eyeball" test is not sufficient to determine the causal impact of the merger on prices, it does highlight that there is no large, obvious impact, and that methodological choices are likely to have a meaningful impact on the findings.

Three of the four mergers occurred during unique periods where overall economic trends are first-order determinants of price and passenger counts. Delta received regulatory approval to acquire Northwest in the fourth quarter of 2008, just as the financial crisis began. This can be seen in the first panel of Figure 1, where prices for overlap and both control groups begin to decrease significantly just after the merger is consummated. The United/Continental merger was consummated in the third quarter of 2010, while the US economy was recovering from the financial crisis. As a result, the pre-merger period includes a substantial decrease in prices, followed by an upward trend that passes through the pre and post-treatment period. The Southwest/AirTran merger also occurred during the economic recovery from the financial crisis, whereas American/USAir coincided with a more stable economic climate. Given the relatively small number of treated units in each merger, there is some concern that unobserved heterogeneous effects of the crisis may bias the results. To the extent that these unobservables are captured in the pre-merger prices, and uncorrelated with treatment, synthetic control and nearest neighbor matching may be well-suited to address the problem. On the other hand, DID using route and time fixed-effects may not be able to address these heterogeneous effects.

Another potential concern is that some routes, like those primarily catering to leisure travelers, are more prone to seasonal fluctuations than others. In Figure 4, passenger counts on the United/Continental overlap routes appear to be highly seasonal, whereas the control routes appear less prone to seasonal demand fluctuations. To a lesser extent, this appears to be an issue in the other three mergers. Time fixed effects will not completely account for this difference between groups, but synthetic control and nearest neighbor could potentially place more weight on routes that exhibit similar seasonal patterns. However, this is only true to the extent that seasonal fluctuations in demand are incorporated into final prices.

In the following section, we show that synthetic control and nearest neighbor tend to produce a better pre-merger fit between the treatment and control group in terms of pre-merger prices. Furthermore, synthetic control and nearest neighbor routes exhibit seasonal demand patterns that are closer to the treatment routes, but are far from a perfect fit. We provide more details in the following subsection. However, as we develop in the final section, synthetic control and nearest neighbor may still suffer from unobserved variable bias.

4.4 Summary of Pre-merger Trends and Matching

In Figure 1 it is not clear, to the eye, whether or not prices follow a parallel trend prior to any of the mergers. As such, we follow Autor (2003) and use regression analysis to perform a formal test. Specifically, we estimate equation (1) but include an interaction between the treatment indicator and a dummy variable for each time period. If these interactions are statistically significant in the pre-merger period then it demonstrates that overlap routes may follow a different time trend prior to the merger, and therefore post-merger effects should be interpreted with caution. We find that in 3 of the 4 mergers, the parallel trend assumption is rejected with 95% confidence. In Delta/Northwest, United/Continental, and Southwest/AirTran, at least 4 of 8 pre-merger interactions are statistically significant with at least 95% confidence when control group 0 is used. Nearly identical results hold when the control group 1 is used, although United/Continental is only statistically significant in 3 of 8 periods with 95% confidence. For American/USAir, only one time period interaction is significant pre-merger with control group 0, and none are significant using control group 1.[24] At best, these results suggest that in all but the American/USAir merger, overlap routes follow different seasonal patterns than do control routes that would be washed out by averaging together quarters, which is apparent in the passenger counts in Figure 4. At worst, the results demonstrate a violation of the parallel trend assumption, and the DID regressions should be interpreted with caution. Nonetheless, we present the DID results in the following subsection, as it is the method most commonly employed in merger retrospectives, including in the airline industry, and serves as an important point of comparison to synthetic control and nearest neighbor matching.

Using the synthetic control method, pre-merger prices on overlap routes look nearly identical to the synthetic control routes. Figure 2 depicts the mean price on overlap and synthetic control routes, estimated using the full set of pre-merger prices. Table 3 summarizes prices and other characteristics on overlap routes and the synthetic control routes.[25] For each

[24]These results are available upon request.

[25]Note that the pre and post merger prices in Tables 2 and 3 do not perfectly match. This is because

of the mergers, the pre-merger fit is very good; the synthetic control routes are able to match the pre-merger seasonality and price trends in all four mergers. Additionally, the synthetic control routes appear to match the post-merger seasonal patterns, although the gap in price levels appear to grow in all but the Southwest/AirTran merger.

While the synthetic controls estimated using the full set of pre-merger prices provide a good in-sample fit, they are never the optimal specification selected using the cross-validation method detailed in Appendix A.1.1. Table 5 summarizes the RMSEs of the backward projections, and in each merger there is a specification with additional competition covariates that produces a lower RMSE. Thus, even though the common practice of using the full set of pre-treatment outcomes creates synthetic control routes that perfectly match the treated routes, adding additional covariates and a subset of pre-treatment outcomes tends to produce a superior out-of-sample prediction.[26]

Although the synthetic control method places positive weight on almost every control route, a small handful of routes receive a large majority of the weight.[27] When using only the full set of pre-merger prices, across all synthetic control routes, the five routes receiving the most weight account for, on average, 85.5% of the total weight.[28] When competition covariates are used with the "optimal" set of pre-merger prices, the top five routes tend to get less cumulative weight: 69%, on average, for both control groups 0 and 1. The relevance of the competitive controls in predicting pre-merger price trends, however, is minimal; on average across mergers, HHI, the number of legacy carriers, and the share of legacy carriers receive weights of .027, .013, and .018, respectively. Moreover, the competition convariates receive no weight in more than 75% of the synthetic control routes. In turn, including these additional matching variables has little impact on the amount of weight placed on an individual route in constructing the synthetic control. Comparing the "full set of pre-merger prices" and the "optimal" specifications, we find that the average absolute difference in the weight placed on a control route is 0.003 across all mergers and synthetic control routes.[29] Thus, it appears that the improvement in prediction of the optimal specification is largely driven by two factors: (i) using a subset of prices may protect against model overfit and (ii) competition covariates,

prices are matched in synthetic control using logs. We transform these values to levels in Table 3 for ease of interpretation. Table 2 reports raw averages.

[26]For Delta/Northwest, the optimal specification includes the full set of pre-merger prices in addition to the competition covariates.

[27]The small positive weights received by most control routes may be an issue with computer precision.

[28]This statistic is calculated for control group 0. For control group 1, the average cumulative weight for the top five routes is 83.5%.

[29]This statistics is calculated for control group 0. For control group 1, the average absolute difference is 0.006.

while having a de minimis effect on the weight an individual route receives, may have larger effects when the impact is aggregated over hundreds of control routes.

In nearest neighbor matching, the parallel trends assumption clearly holds for at least three of the four mergers. As shown in Figure 3, the Delta/Northwest and United/Continental mergers have virtually identical treated and control pre-merger price trends. Lastly, the difference in levels for American/USAir is minimal, especially between the treated group and control group 1. Control group 0 remains relatively close to the treated route throughout the entirety of the pre-merger period and follows a similar, though not identical, seasonal trend.

Interestingly, while synthetic control and nearest neighbor both satisfy the parallel trend test, they tend not to place weight on the same control routes. Across mergers, only 5%-16% of the five "nearest neighbors" were also the top five most heavily weighted routes by the synthetic control method. Both methodologies are intended to remove researcher intuition from the process of selecting control groups, and allow the data to speak for itself. Yet, these two methods result in almost completely dissimilar control groups, both in weighting and composition.

Finally, we note that even though synthetic control and nearest neighbor satisfy the parallel trends test, it does not guarantee that they do not suffer from unobservable variable bias. In Appendix B, we develop a test for whether the treated and synthetic (or nearest neighbor) routes produce prices from the same equilibrium supply and demand conditions. We find that in all mergers, neither synthetic control nor nearest neighbor satisfy these conditions. We therefore show in these Appendix results that, even though both methods satisfy the parallel trend test, they may be just as likely as OLS DID to suffer from unobservable variable bias.

5 Merger Retrospective Results

We now present the merger retrospective results. The focus is on the price impact of the merger, but we make reference to the total passenger results presented in a web appendix. We begin by analyzing each merger separately. A summary of all of the results are presented in Table 6.[30]

We include two years of quarterly, route-level data before and after each merger, and

[30]We perform a number of robustness checks that are reported in a web appendix available on the authors' websites. These include, for each method, estimates that exclude monopoly routes from the sample. It also includes specifications that drop treated routes that experienced post-merger entry or exit. Results are qualitatively the same as those reported in this article.

drop the quarter in which the merger occurred. This results in 16 observations for each route included in the regression.

The DID estimates include two specifications,

$$y_{mt} = \alpha + \beta V_m Post_t + \gamma_m + \lambda_t + \epsilon_{mt} \tag{5}$$

$$y_{mt} = \alpha + \beta_1 V_m Post_{1t} + \beta_2 V_m Post_{2t} + \gamma_m + \lambda_t + \epsilon_{mt}. \tag{6}$$

For the primary results, y_{mt} is the log of the average quarterly price on each route. In equation (5), $V_m Post_t$ is the interaction of an overlap indicator and an indicator for the eight post-merger quarters. In equation (6), we separately estimate treatment effects for the first and second post-merger years.

For the synthetic control method, we include two specifications. The first matches only on the eight pre-merger price observations, and the second includes the three competition covariates and the optimal set of pre-merger prices. In some specifications, the synthetic control estimation routine does not converge for all overlap routes. We exclude overlap routes that did not converge when estimating the treatment effect. For nearest neighbor matching we estimate two specifications: matching on only the eight pre-merger prices, and matching on the eight pre-merger prices and three competition covariates.

5.1 Delta/Northwest results

The Delta/Northwest merger was the first in a wave of airline mergers that began after years of eroding industry profits following the 9/11 terrorist attacks. The Delta/Northwest merger coincided with the onset of the 2008 financial crisis, which impacted many sectors of the economy including the airline industry. Given the small number of overlap routes (5), there is some concern that any estimated effect could be due to unobservable, heterogeneous impacts of the financial crisis. Therefore, synthetic control and nearest-neighbor may be better equipped to produce reliable estimates, as they more closely match pre-merger price trends leading up to the economic downturn.

Table 7 reports the DID estimates, and there is some support for prices decreasing in the first post-merger year, relative to routes where no merging firm was present pre-merger. There is some evidence that overlap route prices increased relative to routes with one merging firm pre-merger. This suggests that the merger resulted in efficiencies but, due to increased market power, pass-through was lower on overlap routes than routes with only one merging firm. The effect on total passengers is inconclusive relative to control group 0, and significantly negative relative to control group 1. Thus, DID offers some support for merger-specific efficiencies lowering overall prices.

On the other hand, synthetic control yields generally insignificant results. Whether or not competition variables are used in addition to pre-merger prices, there is no statistically significant difference between overlap and the synthetic control routes constructed from either of the two control groups.[31] Nearest neighbor matching also produces insignificant results relative to the first control group, whether or not competition controls are included. On the other hand, relative to routes with one merging firm, prices increase significantly by 7%-14%, depending upon the post-merger quarter. These results largely lose significance, however, when extra competition matching variables are included with the optimal set of pre-merger prices.

In total, the three methodologies produce disparate results. Standard DID suggests that there may have been cost efficiencies that led to lower price in the first post-merger year. However, merger-specific cost efficiencies would typically persist into the second post-merger year, which does not seem to be the case. Furthermore, neither synthetic control nor nearest neighbor appear to support these findings. As this merger occurred during a volatile time in the US economy and the parallel trend assumption does not hold in standard DID, we place more weight on synthetic control and nearest neighbor matching. However, as the pre-merger period largely occurred prior to the financial crisis, it is not clear how well pre-merger price trends predict the post-merger, unobserved impact of the crisis. Given the disparity across methodologies, we believe the impact of merger on prices is unclear from the analysis.

5.2 United/Continental results

The United/Continental merger was the second legacy airline combination during the merger wave. It was completed in late 2010, two years after Delta/Northwest received antitrust approval. The two airlines had seven overlap routes prior to the merger, all of which contained endpoints at Newark International, Chicago O'Hare, or Denver International. Unlike the Delta/Northwest merger, all three methodologies tell the same general story; the United/Continental merger resulted in higher prices and no cost efficiencies.

The DID estimates presented in Table 10 show a 7% price increase in the first post-merger year relative to both control groups. The nearly identical results across control groups

[31] The one exception is for the third and eighth post-merger quarters, where prices are 9.4% and 1.4% lower with 95% confidence, when pre-merger prices and competition controls are used for matching and with control group 0.

is consistent with news reports[32] of the two airlines having difficulty combining operations, and therefore being unable to realize marginal cost savings. The sign and magnitude of the estimated effects in the synthetic control and nearest neighbor methods (Tables 11 and 12, respectively) are consistent with the DID findings. The estimated effects across control groups are in the 0%-14% range, depending upon the post-merger quarter. Statistical significance, however, is not always robust to whether or not we match on competition variables in addition to pre-merger prices. Furthermore, while passengers are estimated to fall on overlap routes in each of the three methods, the decrease is not statistically different from zero in any specification other than one version of nearest neighbor (control group 0) and one quarter of one specification of synthetic control (control group 1).

In total, all methods tell a story of increased prices and no marginal cost efficiencies from the merger. Statistical significance, however, is somewhat sensitive to including competition matching variables.

5.3 Southwest/AirTran results

The Southwest/AirTran merger is the only combination of two low-cost carriers in the analysis. Prior to the merger, Southwest and AirTran overlapped on 13 non-stop routes, which is the most overlaps of the four mergers. 12 of the 13 routes include an endpoint at either Baltimore-Washington or Orlando International airports. Previous research has demonstrated that low-cost carriers have a unique competitive constraint on prices (Brueckner, Lee, and Singer (2013)). Furthermore, both airlines are thought to have lower cost structures than legacy airlines, which may limit the scope for merger-specific marginal cost efficiencies. In total, these factors may increase the potential for anti-competitive harm on overlap routes. The results, however, are decidedly mixed.

In Table 13, the DID results strongly indicate a post-merger exercise of market power on overlap roots. Prices increased on overlap routes by about 9%, and this holds even when the effect is separately estimated for the two post-merger years. The estimates are one to three percentage points lower when the control group includes only routes where one merging firm was present pre-merger. This suggests that there were not significant marginal cost efficiencies and that prices increased relative to both control group due to the exercise of market power. Across the 4 specifications, price effects are estimated precisely. In Table 30, significant passenger decreases are also estimated across the four specifications. The standard DID methodology, therefore, paints a clear picture of post-merger exercise of market power

[32] See, for example, Carey and Nicas (2015).

that lead to higher prices and lower quantity.

The other methodologies, however, offer a different story of how the merger affected overlap markets. The synthetic control method estimates do not reject the null hypothesis with 95% confidence when estimating the effect relative to control group 0. Using control group 1 and only matching on pre-merger prices, the merger is estimated to significantly increase price by 2%-3%, depending on the quarter. The estimates are more precise when competition variables are used in addition to pre-merger prices to construct the synthetic control routes. In total, the synthetic control yields supportive evidence of the merger increasing prices, but by a lesser magnitude than when using OLS DID.

In contrast, nearest-neighbor matching yields large and significant negative effects when matching only on pre-merger prices and using control routes where Southwest and AirTran had no pre-merger presence. In the first post-merger year, the price reductions are between 9% and 19% and significant with 99% confidence in the first post-merger year. However, these results flip, or become insignificant, by adding three additional competition matching variables; a positive and significant price effect of 18% is estimated in the third quarter post-merger. All other results in the first year are insignificant. In the second year after the merger, we estimate a significant and negative price effect between 13% and 19%. While adding these additional matching criteria to the synthetic control had a minimal impact, they swing the results greatly in nearest neighbor matching. Using control group 1, the results are more consistent; post-merger prices increased. Still, statistical significance with 95% confidence is dependent upon adding competition matching variables.

Nearest neighbor produces a relatively poor pre-treatment fit in this merger, as can be seen in Figure 3, which may lead to biased estimates. Ferman and Pinto (2021) find that synthetic control is generally biased if pre-treatment fit is imperfect and treatment assignment is correlated with unobserved confounders. Abadie and L'Hour (2021) demonstrate that nearest neighbor and synthetic control are of the same class of estimators. Therefore, nearest neighbor may be subject to the same issues derived in Ferman and Pinto (2021). As mentioned above, the Southwest/AirTran merger (as well as the others) fail to pass our test of unobserved variable bias in nearest neighbor matching. Taken as a whole, these findings show that the nearest neighbor estimates in this merger are likely biased, and little weight should be placed on them.

In summary, the estimated impact of the Southwest/AirTran merger is highly dependent upon both the methodology and the variables used for matching treatment to control routes, especially for the control group without pre-merger overlap. This highlights that reasonable, yet different choices made by the researcher can yield contrasting policy implications.

5.4 American/USAir results

American Airlines and USAir is the most recent legacy airline merger in the United States, and was completed in November of 2013. The merger left the United States with only three legacy airlines. The US Department of Justice initially sued to block the merger, but eventually reached a settlement. As part of the settlement, the merged company had to sell slots and/or gates at six airports. We exclude routes that experienced entry due to the divestiture from the analysis. The merger resulted in 10 overlap routes that did not experience post-merger entry due to divestiture. The overlaps predominantly had endpoints at the airlines' major hubs in Charlotte, Philadelphia, and Dallas.

Similar to the United/Continental merger, all retrospective methodologies point in the same direction. However, in this case, the merger appears to have resulted in a price decrease. Table 16 presents the DID results and, relative to both control groups, overlap prices fell by 7%-10%. Between the two control groups, there is no statistical difference in the magnitude of the price decreases. This suggests that marginal cost efficiencies were realized exclusively on overlap routes. Alternatively, overlap routes received an unobserved negative demand shock subsequent to the merger. However, this is unlikely, as total passengers are also estimated to increase post-merger on overlap routes (Table 33).

Synthetic control results are reported in Table 17 and tell a story similar to DID, although magnitudes are smaller and statistical significance is not as robust. When control routes do not include either merging firm, post-merger price effects are in the -2% to -9% range for all but one quarter, and they are only statistically different from zero in the first and eighth post-merger quarters. Using competition covariates and the optimal subset of prices flips the sign for five of the eight post-merger quarters. Relative to control group 1, overlap prices are found to decrease significantly using both sets of matching variables.

The nearest neighbor estimates are reported in Table 18. Using the control group with no merging firms present, the effect of the merger is estimated to be negative and significant with at least 90% confidence in 7 of 8 post-merger quarters. However, including the competition measures in the matching criteria flips the sign of the estimates in 3 of 8 quarters and statistical significance is lost in all but the first two post-merger quarters. Estimates are more robust when matching on the control group with one merging party present pre-merger. The merger is estimated to lower prices in every quarter, whether or not additional competition matching variables are used. The effects are also statistically different from zero in most post-merger quarters.

The passenger results reported in Appendix Tables 33 through 35 largely agree on a statically significant increase in post-merger passengers on overlap routes. Taken as a whole, the results are consistent with the merger leading to lower prices on overlap routes. However, the findings are not entirely robust to adding competition matching criteria in the synthetic control method. As this is the only merger that passed the parallel trend test, the standard DID results may be more reliable than in the other three mergers.

6 Merger Impact on Control Routes

The results in the previous section demonstrate that the average treatment effect of mergers on overlap routes is not uniform in sign or magnitude. This may be due to heterogeneous realizations of cost efficiencies and/or the ability to exercise market power. Furthermore, it may be difficult to tease out an effect from the data due to there being a small number of overlaps in each merger. In the following two sections, we demonstrate that it may still be possible to draw general conclusions about the impact of airline mergers. First, we estimate the extent to which "control" routes may be impacted by the merger. Not only would this lead to biased estimates of treatment on overlap routes, but, more importantly, should be an additional concern for antitrust enforcers. Second, we analyze if the effect on overlap routes is predictable at the individual route-level, even if the average effect across routes is unclear. This more narrowly focused analysis may potentially provide guidance to antitrust practitioners when investigating future airline mergers.

Merger retrospective typically assume that anticompetitive effects only occur in markets where both firms compete. However, there is a mechanism by which markets with only one firm present pre-merger can be impacted by a merger. In particular, if multi-market contact facilitates collusion in an industry, then a merger can lead to increased coordination in markets where only one merging firm is present pre-merger. Ciliberto and Williams (2014) and Ciliberto, Watkins, and Williams (2019) both find that multi-market contact facilitates tacit collusion and leads to higher prices in the airline industry.

We exploit the fact that a merger leads to an exogenous increase in multi-market contact in markets where only one merging firm is present pre-merger. This increase in multi-market contact will increase the incentive and ability to coordinate prices on these routes, and therefore may lead to higher prices. To begin the analysis, we employ the definition of multi-market contact developed in Ciliberto and Williams (2014). First, we define mmc_{kh}^{t} as the total number of routes on which both airlines k and h offer non-stop service at time t.[33]

[33]As in the previous analysis, we define a firm as offering non-stop service if their market share is greater than 10% in a given quarter.

Then, to measure multi-market contact at the route-level, we define the average amount of multi-market contact across all airlines on a route:

$$AvgMMC_{mt} = \frac{1}{F_{mt}(F_{mt}-1)} \sum_{k=1}^{F} \sum_{h=1,h\neq k}^{F} \mathbf{1}[\text{k and h active}]_{mt} \cdot mmc_{kh}^{t}. \tag{7}$$

Here, F_{mt} is the number of airlines operating non-stop service on route m at time t. $AvgMMC_{mt}$ therefore measures the average amount of multi-market contact across all pairs of airlines on route m at time t. Note that $AvgMMC$ is undefined for monopoly routes, of which there are many in the airline industry. We therefore drop monopoly routes from the analysis.

A merger will lead to a discrete jump in $AvgMMC$ for markets in which one merging firm competes. To see this, by way of example, suppose AA and DL are the only two airlines operating on route m. Also, suppose that mmc between AA and DL is 20, and that mmc between DL and US is 10. If AA and US merge then, all else equal, mmc between AA and DL increases to 30, as all US routes become AA routes. In turn, the incentive for AA and DL to coordinate prices on route m increases, despite the fact that market structure on the route is unchanged.

Table 19 summarizes multi-market contact on routes with only one firm present prior to the merger, excluding monopoly routes. In each of the four mergers, these routes experienced a significant increase in multi-market contact. Interestingly, each successive merger sees a larger increase in multi-market contact. This demonstrates an increasing trend in consolidation across the airline industry and that the incentive for airlines to coordinate has steadily increased over time. Figure 7 depicts average multi-market contact on these same routes for the 8 pre and post-merger quarters. In all four mergers, the discrete jump in multi-market contact is apparent following the merger. The size of the increase following each merger is orders of magnitude larger than any other quarter-to-quarter change in multi-market contact. Only in the Delta/Northwest merger is the increase relatively small, but, as reported in Table 19 it is a statistically significant increase. In the following analysis, we exploit this plausibly exogenous (at the route-level) increase in multi-market contact to estimate whether it leads to an increase in prices.

To analyze whether prices increased on control group 1 routes due to the mergers, we first calculate the average value of $AvgMMC_{mt}$ in the year prior to the merger and subtract it from the average value in the year following the merger. We use the change in the yearly average value of $AvgMMC_{mt}$ in order to control for seasonality. For each route, we denote this change in average multi-market contact as $\Delta AvgMMC_m$. We then estimate a standard DID regression of the following form:

$$y_{mt} = \alpha + \beta D_t \cdot \Delta AvgMMC_m + \gamma_m + \lambda_t + \epsilon_{mt}. \quad (8)$$

Here, y_{mt} is the route-level log price, D_t indicates post-merger time periods, γ_m is a route fixed-effect, and λ_t is a time fixed effect. In this regression, β estimates the average effect of the merger-induced change in multi-market contact on route-level prices.

We estimate two versions of equation (8). The first only includes non-monopoly routes where one merging firm operated prior to the merger. In this specification, all routes receive "treatment," and β captures whether routes that received greater treatment have greater post-merger prices. In the second version, we also include routes where no merging firm was present prior to the merger. In that specification, routes where no merging firm was present prior to the merger serve as a "control" against which to measure the treatment effect. While these routes experience some change in multi-market contact, due to entry and exit across routes, it should be unrelated to the merger.

Table 20 presents the results of both sets of regressions for each of the four mergers. The odd numbered columns are the specifications that only include routes with one merging firm; the even numbered columns are the specifications with both types of routes. In all mergers, there is evidence that an increase in multi-market contact caused an increase in prices. In odd-numbered specifications, the effect is always estimated to be positive and significant with at least 90% confidence. In the specifications with both types of routes, 3 of 4 specifications are positive and significant with at least 95% confidence. Within each merger, other than Delta/Northwest, the estimates are qualitatively the same across the two specifications. In Delta/Northwest, one specification is estimated to be negative, but not different from zero. This merger experienced the smallest average change in multi-market contact, which may explain the discrepancy.

Overall, the evidence strongly indicates that the increase in multi-market contact generated from the mergers caused prices to increase on routes where one merging firm operated prior to the merger. The coefficient estimates translate into an approximately $1 increase in airfare for each merger.[34] While these estimates appear relatively small in magnitude relative to average airfare on these routes ($145-$190), a few points are worth noting. First, when estimating the model with prices in levels, the estimated effect increases to $2 to $7, depending on the merger. So, while the sign of the results is not sensitive to

[34]This excludes the second specification in the DL/NW merger.

levels vs. logs, the magnitudes are much larger when estimating the model in levels. Second, there are approximately 10 times as many routes that experience a change in multi-market contact than there are overlap routes in airline mergers. As such, any harm is aggregated over a far greater number of markets and consumers.

These findings have important consequences for merger retrospectives and antitrust investigations. For merger retrospectives, it demonstrates that "control" routes may be affected by mergers if coordination via multi-market contact is a concern. This, in turn, can bias the estimates of the effect in markets where merging firms compete head-to-head. While we show that routes with one merging firm are impacted by the merger, it is possible that routes with no merging firms are impacted by the merger. This would occur if non-merging firms have an increased incentive to coordinate, due to spill-over effects from coordination on other routes. For example, suppose two non-merging firms interact in a market with a merging party and coordination begins post-merger. This could increase the incentive for the non-merging firms to coordinate in markets with no merging firms. While we cannot directly test this hypothesis, it is another channel through which merger retrospective results may be biased downward. More importantly, it represents more harm to consumers resulting from the merger and warrants further consideration from antitrust authorities, especially in the airline industry.

7 Individual Route Effects

In this section, we study the extent to which price effects on overlap routes are predictable at the route-level, based upon observable pre-merger characteristics. The aim of this analysis is twofold. First, we seek to understand why different mergers led to different outcomes. Second, we hope to provide guidance to antitrust agencies when forecasting the competitive effect of mergers on overlap, non-stop routes.

To begin, we collect the route-level effects that are estimated when using synthetic control and nearest neighbor to estimate the average treatment effect of the mergers. We then use the route-level estimates as a dependent variable in a Lasso variable selection regression. We focus on the estimates of the net impact of the merger (measured against control group 0). This is likely of most interest to antitrust authorities, and less likely to be prone to the bias discussed in the multi-market contact section.

Figure 8 plots the estimated route-level effects for each merger, separately for synthetic control and nearest neighbor. These plots offer insights into the results above, as well as the ability to predict route-level effects. First, there appears to be a high correlation in the route-level estimates of synthetic control and nearest neighbor; the only exception being a few routes in the Southwest/AirTran merger. Second, even in mergers where the average treatment effect appears significantly different from zero (UA/CO and AA/US), there is still a lot of dispersion at the route level. Finally, all mergers have routes with route-level effects that are both positive and negative.

We now use the data points depicted in Figure 8 as the dependent variable in two Lasso regressions. One challenge in explaining route-level price effects in this context is that there are a small number of affected routes, but a large number of potential explanatory variables. Lasso, a machine learning technique, selects only the most relevant explanatory variables by penalizing the number of independent variables in the model. In particular, we implement a leave-one-out cross-validated Lasso routine, which uses a number of Lasso estimators to find the optimal value for the lambda tuning parameter. Each route is held out for one iteration and a sequence of candidate lambda values are used to fit Lasso regressions. The lambda value that minimizes the average mean square error over the iterations is selected for the final Lasso estimation.[35] We then use this optimal lambda tuning parameter to fit a Lasso model which provides a coefficient estimate of the predictive influence of each independent variable. Since Lasso does not directly provide standard errors or measures of uncertainty, we implement a bootstrapping procedure that estimates 500 Lasso models with randomly sampled data. The observed distribution provides 95% confidence intervals for each parameter estimate.

Table 21 reports the estimation results.[36] Column (1) reports the results of predicting the price effects of synthetic control while column (2) presents the same results for nearest neighbor. First, we note that the HHI, change in HHI, total passengers, and the minimum circuity of a connecting flight[37] are selected as explanatory variables to predict both synthetic control and nearest neighbor effects. Circuity and total passengers both have confidence intervals that exclude 0 in predicting the effects of both models. HHI and the change in HHI have confidence intervals that include 0 for synthetic control. All variables have the same sign across regression specifications. Therefore, the predictive value of pre-merger observables

[35]For more details, see Ahrens, Hansen, and Schaffer (2020)

[36]The Lasso method begins with the following potential explanatory variables: total non-stop passengers, HHI, change in HHI, number of legacy carriers, minimum circuity of connecting flights, total miles of flight, total passengers, mean unemployment at endpoints, mean income at endpoints, mean population at endpoints.

[37]Circuity measures how much farther a passenger has to fly when choosing a connecting flight instead of a non-stop. It is calculated as the distance of the closest connecting flight divided by the distance of the non-stop flight.

does not appear overly sensitive to the method used to predict merger price effects.

The regression results can largely be interpreted in the context of a model where coordination is potentially occurring pre-merger, and even more likely to occur post-merger. This would be consistent with the results in the multi-market contact section, but there are still some findings in Table 21 that are difficult to reconcile under a single model. The pre-merger level of HHI positively affects post-merger price increases, which is consistent with its use by antitrust authorities to screen anti-competitive mergers. However, Nocke and Whinston (2022) note that the change in HHI is a more appropriate metric to predict the unilateral price effect of a merger. Nonetheless, they state that the level of HHI may still be appropriate when coordinated effects are a concern. Still, we find that the change in HHI is negatively associated with post-merger price effects, and it is statistically significant when predicting synthetic control estimates. These findings are difficult to reconcile with any particular model, and may be a consequence of a small sample (34 observations).

Across both specifications, the minimum circuity of a connecting flight is negatively associated with the price effect. The greater is the circuity, the less substitutable a connecting flight is for a non-stop, as it requires flying a farther total distance. Therefore, routes with a greater circuity may be more prone to pre-merger coordination, and therefore experience a smaller post-merger price change. Indeed, in its complaint against the AA/US merger, the DOJ alleged that USAir was using its connecting flights to prevent non-stop coordination among legacy carriers.[38] Finally, we find additional route characteristics are significantly associated with post-merger price changes. The total distance of the flight, unemployment rate, and mean income of the population all have predictive value in at least one specification.

In total, we believe the results to be useful in helping antitrust authorities screen for routes on which anti-competitive effects may be realized post-merger. Even if the average impact of a merger is expected to be close to zero on overlap routes, we show that individual routes may experience the exercise of market power. As such, the results in this section may help guide antitrust authorities as to which routes require divestitures to fix the potential problem.

8 Conclusion

This article makes a number of important contributions to the merger retrospective and airlines literatures. We find that the impact of airline mergers is heterogeneous and may generate

[38] To access the complaint, see `https://www.justice.gov/d9/atr/case-documents/attachments/2013/08/13/299968.pdf`.

either higher or lower prices. This heterogeneity is explained, in part, by different pre-merger characteristics on overlap routes. The estimated effect and statistical significance of the average treatment effect on overlap routes are sensitive to both the employed methodology and choices within a given methodology. We identify important sources of estimation sensitivity. First, estimated treatment effects are subject to researcher degrees of freedom in each of three methods, even those developed to limit assumptions about relevance of the control data. Furthermore, both synthetic control and nearest neighbor may be subject to unobserved variable bias. Finally, if coordination via multi-market is suspected, then the control group may also be impacted by the merger and lead to biased estimates.

The American/USAir merger appears to generate lower prices on routes in which the airlines overlapped before the merger. Meanwhile, the United/Continental merger led to higher prices. There are meaningful differences in estimated effects between methods for the Southwest/Airtrain and Delta/Northwest mergers. We estimate that the Southwest/AirTran merger significantly increased prices (7%-10%) using DID regression, but significantly decreased prices (-9% to -19%) using nearest neighbor matching. Meanwhile, the Delta/Northwest merger appears to significantly decrease prices using DID regression but increase prices, with varying significance, using nearest neighbor. Synthetic control produces mixed results for both of these mergers.

However, the estimated effect of the merger on multi-market contact and price increases is more consistent across mergers. The trend in industry consolidation has led to each successive merger being subject to a greater increase in multi-market contact. In turn, each merger has seen an increase in prices on routes where the incentive to coordinate grows more substantially from the merger. This effect likely trumps any harm that is generated from the exercise of market power on overlap routes.

Finally, this paper provides important context for antitrust enforcement agencies seeking to perform merger retrospectives. We find that the workhorse difference-in-differences models often used for merger retrospectives can produce different results than more sophisticated matching algorithms. However, both standard OLS and matching methods may be sensitive to "under the hood" decisions of the practitioner. Furthermore, DID regression results are often used as benchmark by which to measure the accuracy of merger simulations. We demonstrate that this benchmark may not be a precise accounting of actual merger price effects.

Appendix

A Details on Difference-in-Differences Methodologies

We implement each of the difference-in-differences methodologies at the route, m, quarter, t, level. We define a route to be non-stop flight between two airport pairs, and specify routes to be non-directional. For each merger, we define the eight quarters prior to a the merger being consummated as the "pre-merger" period, the eight quarters following consummation as the "post-merger" period, and we drop the quarter in which the merger was completed. A "treated" route is one which each merging firm had at least 10% share for all 8 pre-merger quarters. The two sets of control routes are defined in the body of the paper. The body of the paper provides the technical details for OLS identification. The following subsections provide more detail on synthetic control and nearest neighbor matching.

A.1 Synthetic Control Estimation

In this article, we implement the synthetic control method by applying the extension of Abadie, Diamond, and Hainmueller (2010) developed in Cavallo, Galiani, Noy, and Pantano (2013) for multiple treated units. Again, we define the pre-merger period at time $t = 1, .., r-1$ and the post-merger period to be $t = r, ..., T$. Let Y^N_{mt} be the price that would be observed on route m at time t if no merger occurred, and let Y^I_{mt} be the price if the merger did occur. The effect in route m of the merger at time $t \geq r$ is then defined as $\alpha_{mt} = Y^I_{mt}$ - Y^N_{mt}. In markets affected by the merger, however, Y^N_{mt} is unobserved for $t \geq r$, and therefore must be estimated.

Let C be the number of routes in the control group, which are unaffected by the merger. Also, let Z_m be a $(1 \times z)$ vector of observed predictors for price on route m. Abadie et al. (2010) show that if a set of weights, $w^*_c \geq 0$, exist such that the following conditions hold then an unbiased predictor of α_{mt} can be estimated,

$$\sum_{c=1}^{C} w_c^* Y_{ct} = Y_{mt}, \forall t \in [1,..,r-1] \quad (9)$$

$$\sum_{c=1}^{C} w_c^* Z_c = Z_m \quad (10)$$

$$\sum_{c=1}^{C} w_c^* = 1 \quad . \quad (11)$$

Given these conditions hold, then the following is an estimator for the treatment effect for each post-merger time period in each market:

$$\widehat{\alpha}_{mt} = Y_{mt}^I - \sum_{1=c}^{C} w_c^* Y_{ct}. \quad (12)$$

There is no assurance that a set of weights, w^*, will exist such that equations (22)-(23) hold exactly, nor is a unique solution guaranteed. Finding the set of weights such that each equation approximately holds is the primary estimation challenge.

We implement the following procedure to compute the control weights.[39] Let T_0 denote the number of pre-merger time periods, and $K = (k_1, ..., k_{T_0})$ be a linear combination of pre-merger prices, such that $\bar{Y_m^K} = \sum_{s=1}^{T_0} k_s Y_{ms}$. Let G be a finite set of linear combinations, $K_1, ..., K_G$, and $X_1 = (Z_1'; \bar{Y_1^{K_1}}, ..., \bar{Y_1^{K_G}})$ be a $(n \times 1)$ vector of pre-merger linear combinations and price predictors on a route affected by the merger, where $n = z + G$. Then, define X_0 as a $(n \times C)$ matrix that contains the same pre-merger variables, and each row of X_0 corresponds to a control route. The optimal weights, W^*, are chosen to minimize the distance between the treated unit and a weighted sum of the control units.

More specifically, the following equation is minimized:

$$\|X_1 - X_0 W\| = \sqrt{(X_1 - X_0 W)' V (X_1 - X_0 W)}. \quad (13)$$

Here, V is a $n \times n$ symmetric, positive definite matrix that weights the relevance of the pre-merger prices and characteristics. The synthetic control method is valid for any symmetric, positive definite V matrix; we solve for the optimal V matrix that minimizes the root mean-squared error defined by equation (13).

[39] We implement the method using the "synth" package in R.

We solve for a synthetic control unit for each route where the merging firms both offer non-stop service prior to the merger. We then use equation (12) to calculate the treatment effect on each route. Then, to calculate the average treatment effect of a merger, in each post-merger time period, we take the average of $\widehat{\alpha}_{mt}$ across treated routes: $\bar{\alpha}^a_t = \sum_m^M \widehat{\alpha}_{mt}$.

To test for significance, we follow the permutation tests developed in Cavallo, et al. (2013). First, we calculate a synthetic control route and "placebo" effect for each control route. Then, if there are M routes affected by the merger, we randomly select M placebo effects and take the average. We take 10,000 averages and then test where the average treatment effect falls in the distribution of placebo averages.[40]

Thus, the p-value for the test of whether the average merger price effect is positive is,

$$\text{p-value} = \frac{\sum_{i=1}^{1000} I(\bar{\alpha}^a_t < \bar{\alpha}^{pl})}{1000}, \tag{14}$$

where $\bar{\alpha}^a_t$ is the average treatment effect on the affected routes and $\bar{\alpha}^{pl}$ is an average of M randomly chosen placebo effects on control routes.

A.1.1 Selecting a Synthetic Control Specification with Additional Covariates

Kaul, Klößner, Pfeifer, and Schieler (2015) demonstrate that the common practice of using the full set of pre-intervention outcomes in conjunction with additional covariates to construct the synthetic control is likely to be inappropriate. In particular, even when the additional covariates are important predictors of the outcome variable, the synthetic control algorithm may place negligible weight on the convariates. This incorrect weighting can lead the model to be overfit and rely too heavily on idiosyncratic variations of pre-intervention outcome variables (Abadie, Diamond, and Hainmueller (2015)). This problem may be particularly acute when there are a large number of control units relative to treated units, as is the case in our analysis. To make our results comparable with previous synthetic control studies, we estimate one specification that only uses the full set of pre-treatment outcome variables to estimate synthetic control weights. However, following Kaul, Klößner, Pfeifer, and Schieler (2015), when adding additional covariates, we use a cross-validation method to select a model with a partial set of pre-treatment outcomes.

To select a synthetic control specification that includes covariates in addition to lagged outcomes, we first specify four models with a different set of variables for constructing the synthetic control. All four models contain the same three covariates: the number of legacy

[40] Cavello suggest taking all possible M unit averages. Given the number of control routes, this is computationally infeasible in our samples.

carriers, the total market share of those legacy carriers, and the HHI on the route.[41] Each model also includes a different set of pre-merger prices: (i) the full set of 8 pre-merger prices, (ii) yearly average prices for the two pre-merger years, (iii) the half-year average prices for the 4 half-years prior to the merger, and (iv) prices for pre-merger quarters t-8, t-6, t-3, and t-1. We then test the performance of each model, separately for each merger, by estimating the synthetic control model for each of the four models and projecting the synthetic control route backwards for pre-merger quarters t-16 through t-9. We then select the model that minimizes RMSE between the synthetic control and treated routes in the backwards projection. The model selected is allowed to vary across mergers, as different time-periods and different airlines may yield prices according to different data-generating processes.

A.2 Nearest neighbor matching

To implement nearest neighbor matching, we follow the methodology developed in Abadie and Imbens (2006, 2011). Again, let Y^N_{mt} be the price on route m at time t if no merger had occured, and let Y^I_{mt} be the observed price on a route where the merger did occur. Then, let $\mathcal{J}_N(m)$ be the N "nearest neighbors" of route m. A route's nearest neighbors are determined by first selecting a set of relevant covariates, $x = \{x_1, ..., x_f\}$. Then, the distance between routes m and a are:

$$\|x_m - x_a\|_S = \sqrt{(x_m - x_a)' S^{-1} (x_m - x_a)}. \tag{15}$$

Here, S is a symmetric, positive-definite matrix that weights the relative importance of the covariates, and x_m and x_a are $(1 \times f)$ vectors of observable characteristics for routes m and a, respectively. In this article, we specify S such that we calculate the Mahalanobis[42] distance between routes, which allows us to standardize covariates to be on comparable scales.[43] As in the synthetic control method, we include lagged values of the outcome variable in the characteristic vector, and in some specifications we include the same additional measures of route-level competition that are included in synthetic control method.

The estimated treatment effect on route m is $Y^I_{mt} - \widehat{Y}^N_{mt}$, where:

$$\widehat{Y}^N_{mt} = \frac{1}{N} \sum_{j \in \mathcal{J}_N(m)} \{Y_{jt} + \widehat{\mu}_t(x_m) - \widehat{\mu}_t(x_j)\}. \tag{16}$$

[41] To measure these variables, we use the average value in the year prior to the merger.

[42] See Mahalanobis (1936) for an exact formula.

[43] Results are nearly identical when using the Euclidean distance metric or weighting by the inverse diagonal sample covariate covariance.

In equation (16), $\widehat{\mu_t}$ is a bias correction term developed in Abadie and Imbens (2011) to account for nearest neighbor estimates being inconsistent when matching occurs with more than one continuous variable. The bias correction term, $\widehat{\mu_t}$, is estimated by regressing the outcome variable on the set of matching covariates, and then using the predicted value for each route. The average treatment effect at time t is then:

$$\widehat{\tau}_t = \frac{1}{M} \sum_{m=1}^{M} Y_{mt}^{I} - \widehat{Y}_{mt}^{N}, \tag{17}$$

where M is the number of routes that are affected by the merger. Standard errors are derived as in Abadie and Imbens (2006, 2011).[44]

B Measuring Unobservable Bias in Synthetic Control and Nearest Neighbor

In this section, we analyze the extent to which the DID-style estimators control for "unobservables," even when observed pre-treatment outcomes are well matched between treatment and control groups. We propose a simple statistic that can be generated when the following conditions are met: (i) data is available on a variable that is simultaneously determined in an equilibrium with the outcome variable of interest and (ii) that variable is not included in the matching process. In such a setting, this second variable is "unobserved" in the matching algorithm. It can therefore be measured, post-estimation, how similar this unobserved variable is between the treated and the (constructed) control group. To measure this, we propose calculating the RMSE of the "unobservable" variable between the treated and control groups.

Standard merger retrospectives meet both of the above conditions. To measure the impact of a merger on price, lagged prices and other relevant controls are typically included in the specification (as we do in this article). Yet, equilibrium prices are an outcome of a supply and demand system, which simultaneously determines price *and* quantity. As we now demonstrate, if the estimator is unbiased, then it implies that the synthetic unit should match both the pre-merger price *and* the "unobserved" quantity of the treated unit. Merger retrospectives, therefore, offer an important setting to analyze how well matching estimators control for unobservables.

In the context of a linear factor model, Abadie, Diamond, and Hainmueller (2010) and

[44] See those respective articles for more detail.

Abadie (2021) develop and present the conditions under which a well matched synthetic control unit may result in a biased estimated treatment effect, using the following equation:

$$Y_{jt}^N = \delta_t + \theta_t Z_j + \lambda_t \mu_j + \epsilon_{jt}. \tag{18}$$

Here, Z_j is a vector of predictors (including lagged outcome variables) of Y_{jt}^N, μ_j is a vector of unobserved predictors, and ϵ_{jt} is a zero-mean error term. Abadie, Diamond, and Hainmueller (2010) find that the bias in the synthetic control estimator is bounded by the ratio of the size of transitory shocks, ϵ_{jt}, and the number of pre-intervention periods. Bias is also more likely as the number of control units grows larger relative to the number of pre-intervention periods. Intuitively, if the synthetic control is able to match the observable predictors, Z_j, then bias will arise only if differences in the values of ϵ_{jt} between the treated unit and control units compensate for differences in the unobservable predictors, μ_j. This is less likely to occur if ϵ_{jt}'s are smaller or there are more pre-intervention periods. Additionally, a large number of control units increases the likelihood of over-fitting, as there are a greater number of ϵ_{jt}'s on which to place positive weight.

In practice, researchers typically interpret a good match between pre-treatment outcomes on treated and synthetic control units as evidence of an unbiased estimate. We now demonstrate that in many contexts an additional test of bias may be implemented.

B.1 A test for unobservable variable bias

In our setting the outcome variable, route-level price, is simultaneously determined by market supply and demand conditions. Motivated by the linear factor model results in Abadie, Diamond, and Hainmueller (2010), we develop a test for the extent to which synthetic routes control for unobservable determinants of equilibrium prices.[45] In its simplest form, route-level supply and demand, Q_t^d and Q_t^s, respectively, are determined by a system of linear equations,[46]

$$Q_t^d = \psi_t^d \xi_j^d + \gamma_t^d + \alpha P_{jt} + u_{jt}$$
$$Q_t^s = \psi_t^s \xi_j^s + \gamma_t^s + \beta P_{jt} + v_{jt}.$$

[45]As developed in Abadie and L'Hour (2021), synthetic control and nearest neighbor are both in the same class of estimators, which they call a penalized synthetic control. Thus, the same test we develop in the context of traditional synthetic control is also valid for nearest neighbor matching.

[46]We develop our results in the context of a supply and demand system. The same logic holds for non-linear systems.

Using these equations, we can solve for an equilibrium market pricing function by setting $Q_t^d = Q_t^s$,

$$P_{jt} = (\frac{1}{\alpha - \beta})(\psi_t^s \xi_j^s - \psi_t^d \xi_j^d + \gamma_t^s - \gamma_t^d + v_{jt} - u_{jt}). \tag{19}$$

Then, with a simple change in variables, and subsetting the route-level supply and demand shifters into observable and unobservable components, we arrive at the following market price equation,

$$P_{jt} = \delta_t + \theta_t Z_j + \lambda_t \mu_j + \epsilon_{jt}. \tag{20}$$

Here, $\delta_t = \frac{1}{\alpha-\beta}(\gamma_t^s - \gamma_t^d)$ and $\epsilon_{jt} = \frac{1}{\alpha-\beta}(v_{jt} - u_{jt})$. The matrices $\theta_t Z_j$ and $\lambda_t \mu_j$ are the observed and unobserved components of $\frac{1}{\alpha-\beta}(\psi_t^s \xi_j^s - \psi_t^d \xi_j^d)$, respectively. Equation (20) takes the same form as the linear factor model presented in Abadie, Diamond, and Hainmueller (2010). Now, however, it can be interpreted as the outcome of a supply and demand system.

Now suppose that the synthetic control method produces a set of weights, $\mathbf{w}$, that generate an exact match between pre-merger prices on the synthetic and treated route. We show that if the synthetic control method properly accounts for unobserved determinants of price then the same weights, $\mathbf{w}$, should produce an exact match between pre-merger quantities on synthetic and treated routes.

Let $t \leq k-1$ be the pre-merger time periods. Defining the treated route as $j = 1$ and the remaining $J - 1$ routes as control routes, let $\mathbf{w}$ satisfy the following conditions:

$$\sum_{j=2}^{J} \mathbf{w}_j P_{jt} = \delta_t + \theta_t \sum_{j=2}^{J} \mathbf{w}_j Z_j + \lambda_t \sum_{j=2}^{J} \mathbf{w}_j \mu_j + \sum_{j=2}^{J} \mathbf{w}_j \epsilon_{jt} \tag{21}$$

$$\sum_{j=2}^{J} \mathbf{w}_j P_{j1} = P_{11}, \ \ldots \ , \sum_{j=2}^{J} \mathbf{w}_j P_{j,k-1} = P_{1,k-1}, \text{ and } , \sum_{j=2}^{J} \mathbf{w}_j Z_j = Z_1. \tag{22}$$

In words, $\mathbf{w}$ is a set of weights applied to the control routes such that they produce an exact match with pre-merger prices and additional explanatory variables on the treated route. We now show that if the parameters in equation (21) are derived from an equilibrium of the supply and demand system above and the conditions in equations (21) and (22) hold, then the following conditions must also hold for all $t \leq k-1$:

$$\begin{aligned} Q_{1t}^s &= \textstyle\sum_{j=2}^{J} \mathbf{w}_j Q_{jt}^s \\ Q_{1t}^d &= \textstyle\sum_{j=2}^{J} \mathbf{w}_j Q_{jt}^d \end{aligned} \tag{23}$$

In equilibrium $Q^s_{1t} = Q^d_{1t}$, and therefore the conditions in (23) imply that $\sum_{j=2}^J \mathbf{w}_j Q^s_{jt} - \sum_{j=2}^J \mathbf{w}_j Q^d_{jt} = 0$. Then, substituting in the parameters of the supply and demand system, $\sum_{j=2}^J \mathbf{w}_j Q^s_{jt} - \sum_{j=2}^J \mathbf{w}_j Q^d_{jt} = 0$ can be rewritten as,

$$\sum_{j=2}^{J} \mathbf{w}_j [\psi^s_t \xi^s_j - \psi^d_t \xi^d_j] + \gamma^s_t - \gamma^d_t + \sum_{j=2}^{J} \mathbf{w}_j [\upsilon_{jt} - u_{jt}] + (\beta - \alpha) \sum_{j=2}^{J} \mathbf{w}_j P_{jt} = 0.$$

Using the same variable definitions as in equation (20), this simplifies to:

$$\frac{1}{\alpha - \beta}[\theta_t \sum_{j=2}^{J} \mathbf{w}_j Z_j + \lambda_t \sum_{j=2}^{J} \mathbf{w}_j \mu_j + \delta_t] + \sum_{j=2}^{J} \mathbf{w}_j \epsilon_{jt} + (\beta - \alpha) \sum_{j=2}^{J} \mathbf{w}_j P_{jt} = 0$$

Finally, solving for price we have:

$$\sum_{j=2}^{J} \mathbf{w}_j P_{jt} = \delta_t + \theta_t \sum_{j=2}^{J} \mathbf{w}_j Z_j + \lambda_t \sum_{j=2}^{J} \mathbf{w}_j \mu_j + \sum_{j=2}^{J} \mathbf{w}_j \epsilon_{jt}. \tag{24}$$

Equation (24) is identical to equation (21). It follows that if the two conditions in (23) do not hold then neither does (21). While equation (21) is not directly observed (and therefore testable), condition (23) is directly observable in the data. We have therefore derived a test for whether or not the synthetic control method controls for unobservable determinants of price, even when pre-treatment and synthetic outcomes perfectly match. Specifically, we use quantity data to test if condition (23) holds. If it does not hold then the only way for pre-merger prices and characteristics to exactly match between treated and synthetic routes (condition (22)), is for the scale of the individual shocks, ϵ_{jt}, to compensate for unobserved determinants of price. Finally, it is important to note that if equation (23) holds, this does not rule out the presence of unobservable variable bias. But, if it does not hold then we can jointly reject the assumed demand system and the matching procedure controlling for all unobservables.[47]

B.2 Implementation and results of the unobservable bias test

Q^s_{1t} and Q^d_{1t} in equation (23) are simply the observed quantities on the treated routes. Similarly, $\sum_{j=2}^J \mathbf{w}_j Q^s_{jt}$ and $\sum_{j=2}^J \mathbf{w}_j Q^d_{jt}$ are the weighted-average quantities on control routes, where $\mathbf{w}$ is the vector of weights solved for when price is the outcome variable. We refer to

[47] If the underlying parameters governing the supply and demand systems on the treated and synthetic control routes do not match then equation (23) will fail, which could also lead to biased estimates.

$\sum_{j=2}^{J} \mathbf{w}_j Q_{jt}$ as the unobserved synthetic quantity (or passengers, in our setting). To test if condition (23) holds, we calculate the root-mean squared error (RMSE) between treated route passengers and unobserved synthetic passengers using the 8 pre-treatment time periods.

If the RMSE is large then we have reason to believe that there are unobservable differences between the treated and synthetic route. On the other hand, if the passengers are similar on both treated and synthetic routes then it is likely that prices and quantities on both routes are being generated by similar supply and demand systems, and unobservable bias and/or over-fitting are less of a concern.

This test applies only for synthetic control and nearest neighbor matching, as these estimators attempt to solve for an exact fit between pre-treatment matching variables. On the other hand, the identifying assumptions for OLS only require parallel pre-treatment trends. Nonetheless, we also present results for OLS as a point of comparison. Figures 4, 5, and 6 depict the unobserved passenger trends for the OLS, synthetic control, and nearest neighbor control groups, respectively.[48] In the figures, there are clear discrepancies between the pre-merger passenger trends of the control and treated units, with respect to both levels and the degree of seasonality. It is not immediately clear which method produces the best-fitting control group in terms of unobserved passenger counts, which highlights the importance of the metric developed in the previous subsection.

Table 22 reports the root mean-squred error (RMSE) between treated and unobserved passengers in the 8 pre-merger quarters, for each of the 4 mergers. To implement the test we, use log-prices and and log-quantities. We do this, as we match on log-prices in the main text of tha paper. We are therefore assuming a log-linear demand system (i.e. log quantities are a linear function of log-prices). Table 22 transforms the RMSE of log-passengers into levels so that they can be benchmarked against Figures 4, 5, and 6. For purposes of comparison with OLS, we present results for the synthetic control and nearest neighbor specifications without additional competition controls. The primary finding in this table is that, for synthetic control and nearest neighbor, the RMSE of "unoboserved" passengers is orders of magnitude larger than the RMSE of the pre-merger synthetic control route matched on pre-merger prices.

When matching on pre-merger prices, the synthetic control route is a near perfect fit in almost every instance, with an RMSE close to zero. Conversely, the RMSE of unobserved passengers is, on average, 2,056 across all mergers and specifications. Nearest neighbor

[48]For OLS and nearest neighbor, unobserved passenger counts are calculated as described above for synthetic control. We first select the weights applied to the control routes when estimating the *price* effect. We then use these weights to construct a weighted average of passengers on the control routes.

matching performs slightly better with unobserved passengers (average RMSE is equal to 2,037) and OLS performs slightly worse (average RMSE of 2,174). The large discrepancy between perfectly matched prices and the large RMSE of unobserved passengers suggests that the matching process is unlikely to be controlling for unobserved determinants of pre-merger prices. Furthermore, the matching process does not provide a meaningfully better fit on unobserved passengers than OLS, which does not constrain pre-merger prices to closely match between treated and control routes.

As another benchmark, we compare the RMSE of unobserved passengers to the actual number of pre-merger passengers on treated routes. The average number of passengers on overlap routes, across mergers, ranges from 4,217 (Delta/Northwest) to 8,765 (Southwest/AirTran). For synthetic control, the RMSE of unobserved passengers ranges from 16% (Southwest/AirTran: $\frac{1409}{8765} = 0.16$) to 56% (United/Continental: $\frac{4123}{7420} = 0.56$) of the average number of passengers on pre-merger overlap routes. The range of RMSE for unobserved passengers on nearest neighbor routes, as a fraction of actual passengers is similar: 29% to 54%.

To the best of our knowledge, when fitting a synthetic control unit, there is no formal definition of an acceptable versus unacceptable fit. Similarly, we do not derive a threshold beyond which the RMSE of an unobservable variable is unacceptably high. However, we do note that the "unobservable" passenger counts on the synthetic control and nearest neighbor routes are very similar to those found by equally weighting all control units, as in OLS. This suggests that, while synthetic control and nearest neighbor routes do a better job at fitting pre-merger prices on overlap routes, this need not imply that they better control for unobserved determinants of equilibrium prices. Similarly, it seems unlikely that the implied supply and demand system generating pricing data on synthetic and nearest neighbor routes approximates those on overlap routes, as they cannot reproduce another variable co-determined in equilibrium – quantity. Overall, the results in Table 22 are consistent with synthetic control and nearest neighbor matching being unable to control for unobserved variables and/or suffering from over-fitting. Thus, even though synthetic control and nearest neighbor (unlike OLS) were able to satisfy the parallel trend assumption, they do not necessarily produce a control group that serves as a better benchmark against which to measure treatment effects.

Finally, in Appendix Tables 23 we compare synthetic control and nearest neighbor matching with only pre-merger prices to the corresponding model using additional competition controls. In the case of synthetic control, we use the model determined by the "optimal" specification selection. Interestingly, the "optimal" model has a higher RMSE of unobserved passengers, on average, than the model that only uses pre-merger prices. The use of controls in nearest neighbor matching provides a moderate improvement, as the RMSE drops by almost 900. Taken together, the inclusion of controls and simple cross-validation may provide some improvement over the simplest models, but the additional benefit is marginal.

Our findings demonstrate the importance of carefully constructing the control group in all DID-style methods. For synthetic control, Abadie (2021) summarizes a number of important considerations researchers should account for when applying the estimation method. The use of finely tuned cross-validation methods, while time consuming, may provide assurance that synthetic control groups are not composed of random realizations of the dependent variable. However, cross-validation does not guarantee that synthetic control or nearest neighbor matching better control for unobserved determinants of the outcome variable, even compared to standard OLS difference-in-differences. We develop an additional statistic that can help determine if unobserved determinants of the outcome variable are controlled accounted for in the control unit(s). Furthermore, as best practice, we suggest implementing synthetic control or a similar matching estimator, in addition to standard DID, when measuring the impact of a consummated merger.

References

Abadie, A. (2021): "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects," *Journal of Economic Literature*, 59(2), 391–425.

Abadie, A., A. Diamond, and J. Hainmueller (2010): "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 105(490), 493–505.

Abadie, A., A. Diamond, and J. Hainmueller (2015): "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, 59(2), 495–510.

Abadie, A., and J. Gardeazabal (2003): "The Economic Cost of Conflict: A Case Study of the Basque Country," *American Economic Review*, 93(1), 112–132.

Abadie, A., and G. W. Imbens (2006): "Large Sample Properties of Matching Estimators for Average Treatment Effects," *Econometrica*, 74, 235–267.

——— (2011): "Bias-Corrected Matching Estimators for Average Treatment Effects," *Journal of Business and Economic Statistics*, 29, 1–11.

Abadie, A., and J. L'Hour (2021): "A Penalized Synthetic Control Estimator for Disaggregated Data," *Journal of the American Statistical Association*, Forthcoming.

Ahrens, A., C. B. Hansen, and M. E. Schaffer (2020): "lassopack: Model selection and prediction with regularized regression in Stata," *The Stata Journal*, 20(1), 176–235.

Allain, M.-L., C. Chambolle, S. Turolla, and S. B. Villas-Boas (2017): "Retail Mergers and Food Prices: Evidence from France," *Journal of Industrial Economics*, 65(3), 469–509.

Ashenfelter, O. C., D. Hosken, and M. C. Weinberg (2014): "Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consumated Mergers," *Journal of Law and Economics*, 57(S3), S67–S100.

——— (2015): "Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry," *RAND Journal of Economics*, 46(2), 328–361.

Autor, D. H. (2003): "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing," *Journal of Labor Economics*, (1), 1–42.

Baker, A. C., D. F. Larcker, and C. C. Wang (2022): "How Much Should We Trust Staggered Difference-in-Differences Estimates?," *Journal of Financial Economics*, 144(2), 370–395.

Bernheim, B. D., and M. D. Whinston (1990): "Multimarket Contact and Collusive Behavior," *The RAND Journal of Economics*, 21(1), 1–26.

Berry, S., and P. Jia (2010): "Tracing the Woes: An Empirical Analysis of the Airline Industry," *American Economic Journal: Microeconomics*, 2(3), 1–43.

Borenstein, S. (1990): "Airline Mergers, Airport Dominance, and Market Power," *American Economic Review*, 80(2), 400–405.

Brueckner, J. K., D. Lee, and E. S. Singer (2013): "Airline Competition and Domestic US Airfares: A Comprehensive Reappraisal," *Economics of Transportation*, 2, 1–17.

Carey, S., and J. Nicas (2015): "United Continental Is Still Shaky Five Years After Merger," *The Wall Street Journal*, July, https://www.wsj.com/articles/united–flights–grounded–due–to–computer–issue–1436361911.

Carlton, D., M. Israel, I. MacSwain, and E. Orlov (2019): "Are Legacy Airline Mergers Pro- or Anti-Competitive: Evidence From Recent U.S. Airline Mergers," *International Journal of Industrial Organization*, 62, 58–95.

Cavallo, E., S. Galiani, I. Noy, and J. Pantano (2013): "Catastrophic Natural Disasters and Economic Growth," *Review of Economics and Statistics*, 95(5), 1549–1561.

Ciliberto, F., and E. Tamer (2009): "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, 77(6), 1791–1828.

Ciliberto, F., E. Watkins, and J. W. Williams (2019): "Collusive Pricing Patterns in the US Airline Industry," *International Journal of Industrial Organization*, 62, 136–157.

Ciliberto, F., and J. W. Williams (2014): "Does Multimarket Contact Facilitate Tacit Collusion? Inference on Conduct Parameters in the Airline Industry," *RAND Journal of Economics*, 45, 764–791.

Dafny, L., M. Duggan, and S. Ramanarayanan (2012): "Does Matching Overcome LaLonde's critique of nonexperimental estimators?," *American Economic Review*, 102(2), 1161–1185.

Das, S. (2019): "Effect of Merger on Market Price and Product Quality: American and US Airways," *Review of Industrial Organization*, 55, 339–374.

Deryugina, T., A. MacKay, and J. Reif (2020): "The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation," *American Economic Journal: Applied Microeconomics*, 12(1), 86–114.

Dix, R., and R. Orzach (2021): "Market Power Spillovers Across Airline Routes," Working paper.

Ferman, B., and C. Pinto (2021): "Synthetic Controls With Imperfect Pretreatment Fit," *Quantitative Economics*, 12, 1197–1221.

Garmon, C. (2017): "The Accuracy of Hospital Merger Screening Methods," *The RAND Journal of Economics*, 48(4), 1068–1102.

Goolsbee, A., and C. Syverson (2008): "How Do Incumbent Airlines Respond to the Threat of Entry? Evidence From the Major Airlines," *Quarterly Journal of Economics*, 123(4), 1611–1633.

Hosken, D. S., L. M. Olson, and L. K. Smith (2018): "Do retail mergers affect competition? Evidence from grocery retailing," *Journal of Economics & Management Strategy*, 27(1), 3–22.

Houde, J.-F. (2012): "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline," *American Economic Review*, 102(5), 2147–2182.

Huschelrath, K., and K. Muller (2013): "The Competitive Effects of Firm Exit: Evidence from the U.S. Airline Industry," *Economics of Transportation*, 2, 72–85.

Imbens, G. W. (2015): "Matching Methods in Practive: Three Examples," *The Journal of Human Resources*, 50(2), 373–419.

Kaul, A., S. Klössner, G. Pfeifer, and M. Schieler (2015): "Synthetic Control Methods: Never Use All Pre-Intervention Outcomes Together With Covariates," MPRA Paper 83790, University Library of Munich, Germany.

Kim, E. H., and V. Singal (1993): "Mergers and Market Power: Evidence from the Airline Industry," *American Economic Review*, 83(3), 549–569.

Lazarev, J., A. Nevo, and R. Town (2021): "What Can We Learn from Merger Retrospectives? Lessons from the Airline Industry," Working paper.

Lewis, M. S. (2020): "On the Absence of Directional Price Discrimination in the U.S. Airline Industry," *Journal of Industrial Economics*, Forthcoming.

List, J., D. L. Millimet, P. G. Fredriksson, and W. W. McHone (2003): "Effects of Environmental Regulations on Manufacturing and Plant Births: Evidence from a Propensity Score Matching Estimator," *The Review of Economics and Statistics*, 85(4), 944–952.

Mahalanobis, P. C. (1936): "On the Generalised Distance in Statistics," *Procedings of the National Institute of Sciences of India*, 2(1), 49–55.

McCabe, M. J. (2002): "Journal Pricing and Mergers: A Portfolio Approach," *American Economic Review*, 92(1), 259–269.

Mehta, A., and N. H. Miller (2012): "Choosing the Appropriate Control Group in Merger Evaluation," *More Pros and Cons of Merger Control*, pp. 189–210.

Morrison, S. A. (1996): "Airline Mergers: A Longer View," *Journal of Transport Economics and Policy*, 30(3), 237–250.

Nocke, V., and M. D. Whinston (2022): "Concentration Thresholds for Horizontal Mergers," *American Economic Review*, 112(6), 1915–1948.

Peters, C. (2006): "Evaluating the Performance of Merger Simulation: Evidence From the U.S. Airline Industry," *The Journal of Law and Economics*, 49(2), 627–649.

Sloczynski, T. (2022): "Interpreting OLS Estimands When Treatment Effects Are Heterogeneous: Smaller Groups Get Larger Weights," *The Review of Economics and Statistics*, 104(3), 501–509.

Weinberg, M. C., and D. Hosken (2013): "Evidence on the Accuracy of Merger Simulations," *Review of Economics and Statistics*, 95(5), 1584–1600.

Werden, G. J., A. S. Joskow, and R. L. Johnson (1991): "The Effects of Mergers on Price and Output," *Managerial and Decision Economics*, 12, 341–352.

Wooldridge, J. (2002): *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA.

C Figures

Figure 1: Prices Before and After Mergers

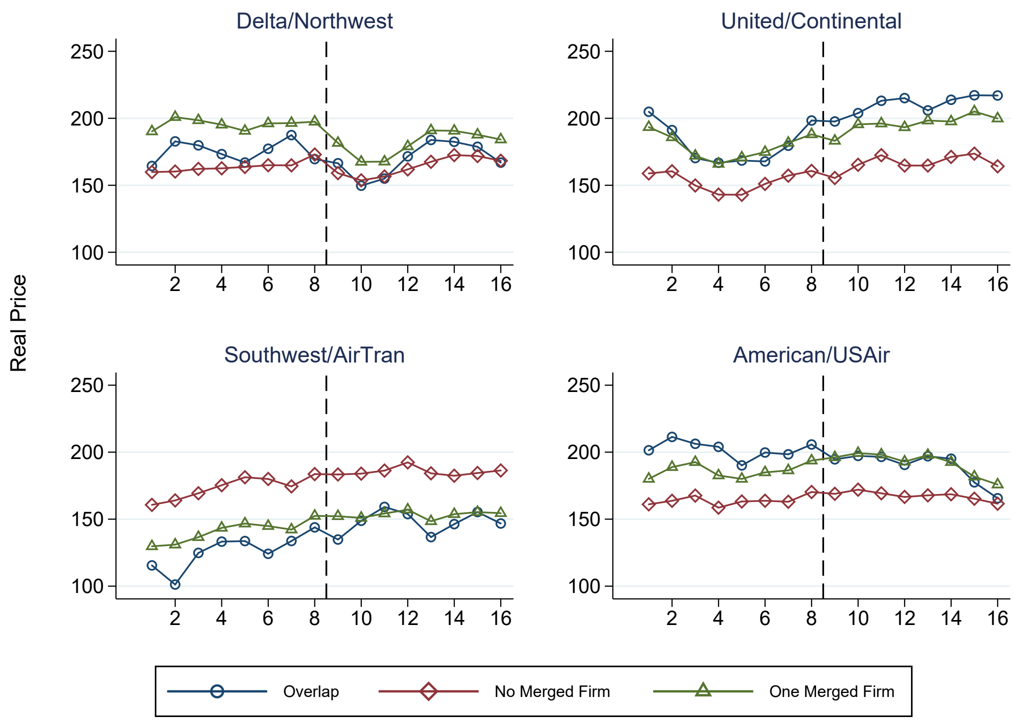

Notes: Each panel depicts quarterly-average, route-level prices for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Overlap indicates routes where both mergers competed prior to the merger. No Merged Firm indicates routes where neither firm were present pre-merger. One Merged Firm indicates routes where exactly one firm competed pre-merger.

Figure 2: Synthetic Control Price Trends

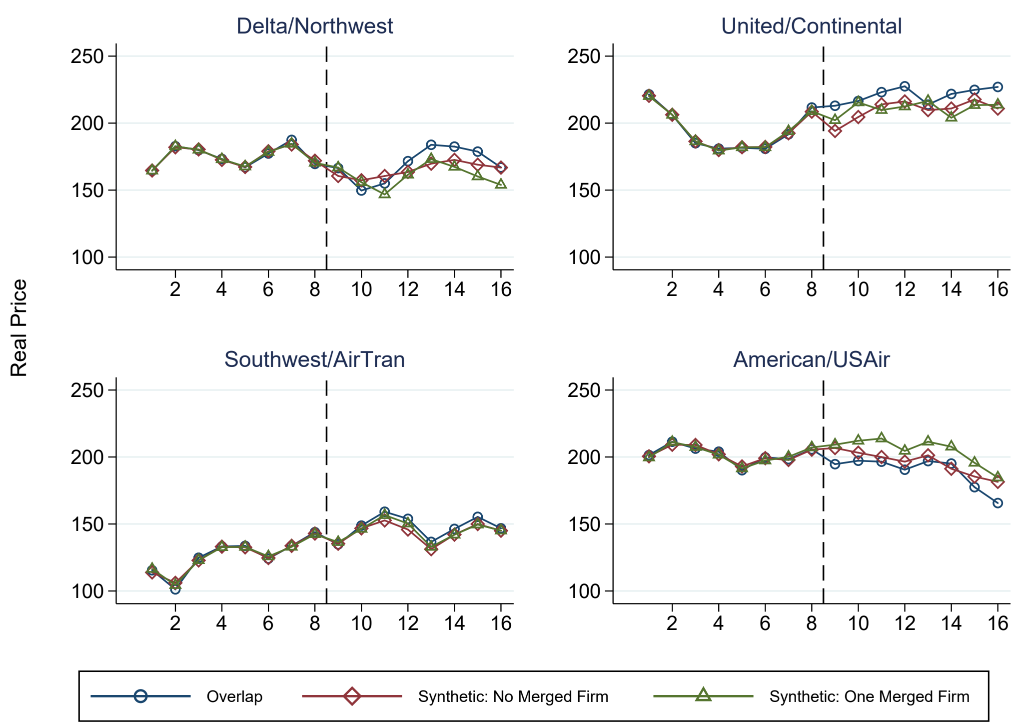

Notes: Each panel depicts quarterly-average, route-level prices for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Overlap indicates observed prices on routes where both mergers competed prior to the merger. No Merged Firm indicates prices on the "synthetic" route estimated using the synthetic control method with only routes where neither firm were present pre-merger. One Merged Firm indicates prices on the "synthetic" route estimated using the synthetic control method with only routes where exactly one firm competed pre-merger.

Figure 3: Nearest Neighbor Price Trends

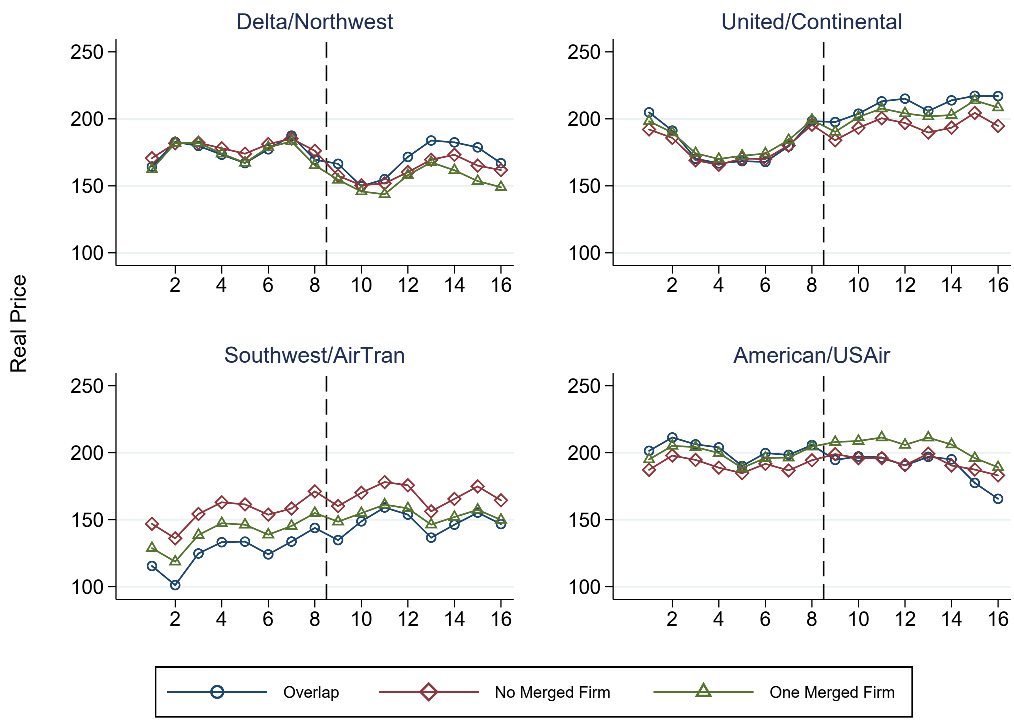

Notes: Each panel depicts quarterly-average, route-level prices for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Overlap indicates observed prices on routes where both mergers competed prior to the merger. No Merged Firm indicates prices on "nearest neighbor" routes estimated using nearest neighbor matching with only routes where neither firm were present pre-merger. One Merged Firm indicates prices on the "nearest neighbor" routes estimated using nearest neighbor matching with only routes where exactly one firm competed pre-merger.

Figure 4: Passengers Before and After Mergers

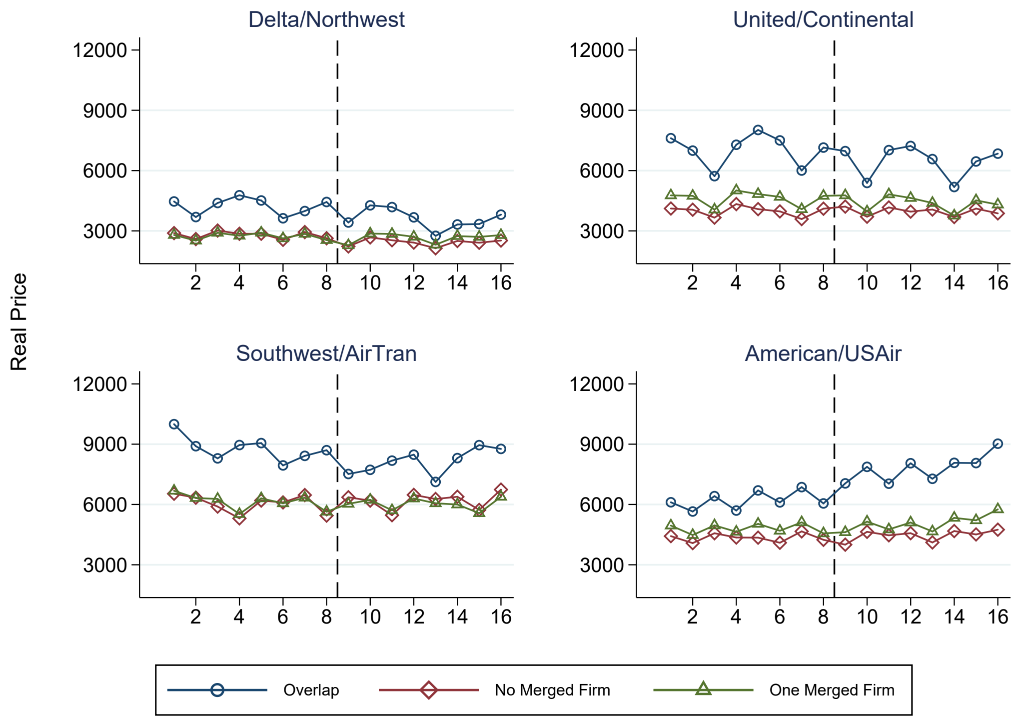

Notes: Each panel depicts the average number of passengers for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Overlap indicates routes where both mergers competed prior to the merger. No Merged Firm indicates routes where neither firm were present pre-merger. One Merged Firm indicates routes where exactly one firm competed pre-merger.

Figure 5: Unobservable Passengers on Synthetic Control Routes

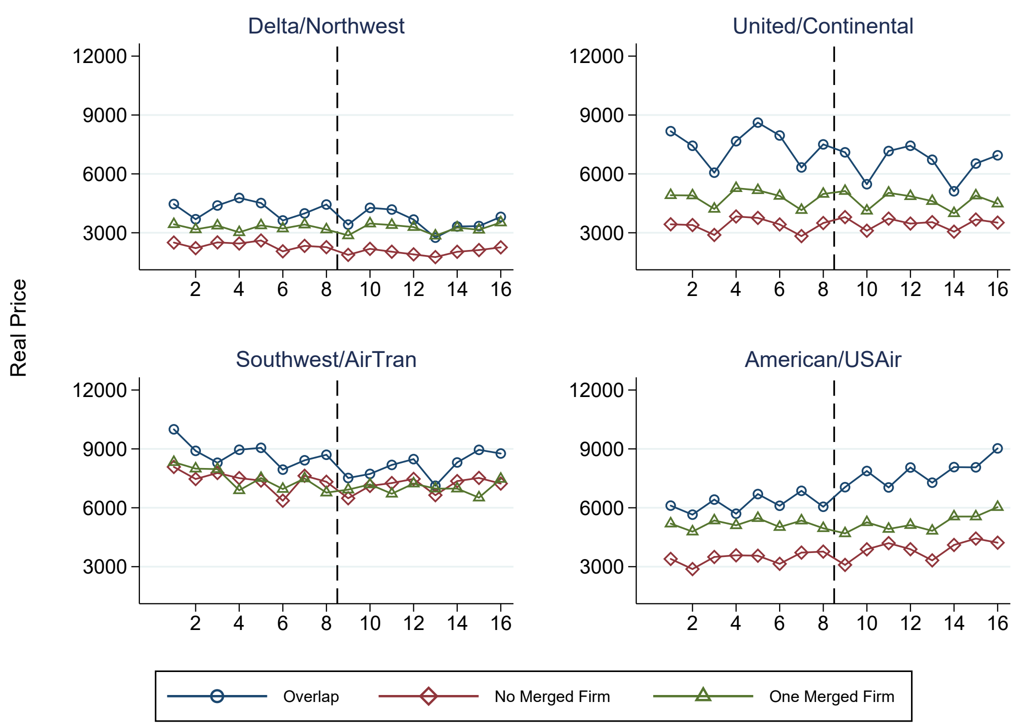

Notes: Each panel depicts the average number of passengers for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Overlap indicates routes where both mergers competed prior to the merger. No Merged Firm and One Merged Firm indicate routes where no merged firm or exactly one merged firm were present, respectively, pre-merger. Passengers on these routes are a weighted average, where the weights are estimated using the synthetic control method, *price* is the outcome variable and only pre-merger prices are used for matching.

Figure 6: Unobservable Passengers on Nearest Neighbor Routes

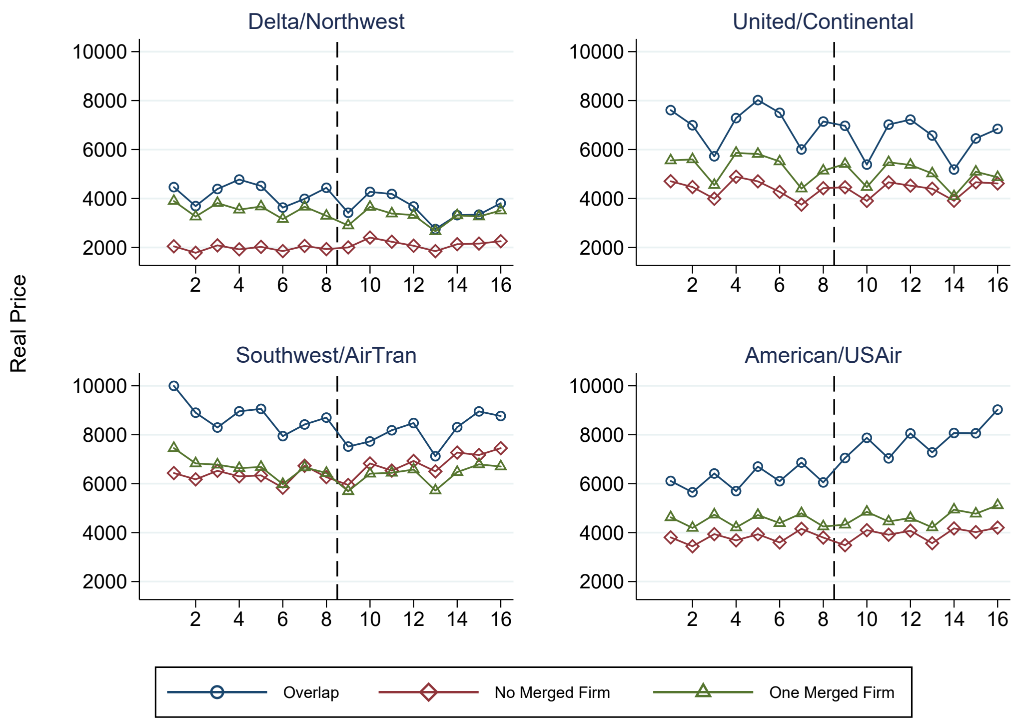

Notes: Each panel depicts the average number of passengers for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Overlap indicates routes where both mergers competed prior to the merger. No Merged Firm and One Merged Firm indicate routes where no merged firm or exactly one merged firm were present, respectively, pre-merger. Passengers on these routes are a weighted average, where the weights were estimated using the nearest neighbor matching, *price* is the outcome variable and only pre-merger prices are used for matching.

Figure 7: Average Multi-Market Contact

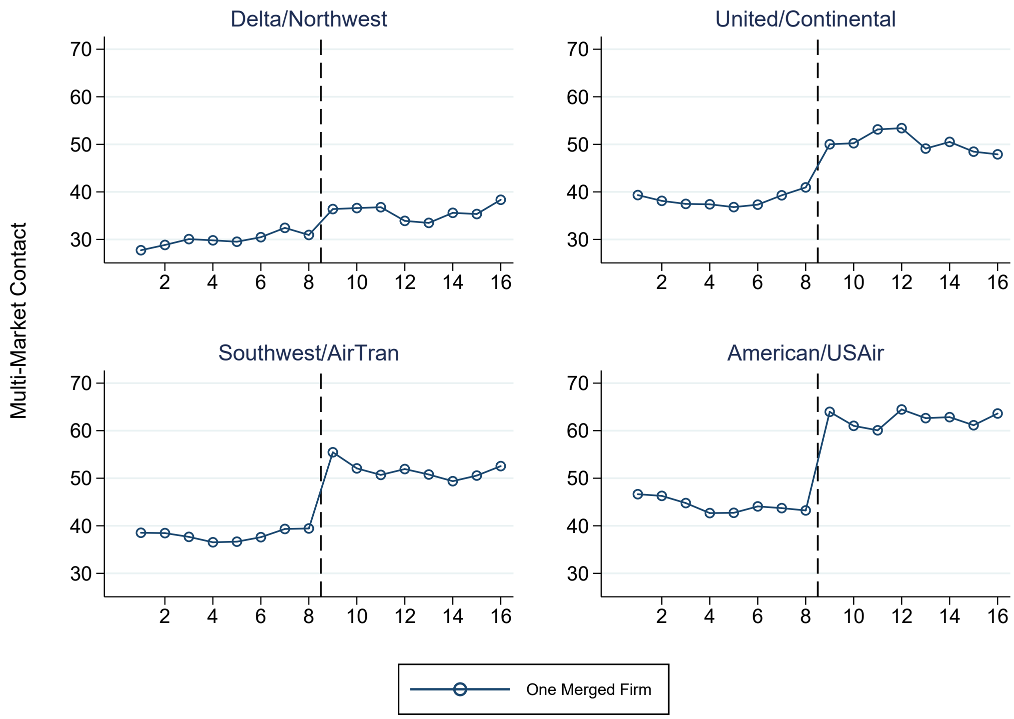

Notes: Each panel depicts the average multi-market contact for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Only routes with one merging firm present pre-merger (control group 0) are used to construct the figure. Only routes with at least two airlines present for all 8 pre and post-merger quarters are included.

Figure 8: Route-level Estimated Price Effects

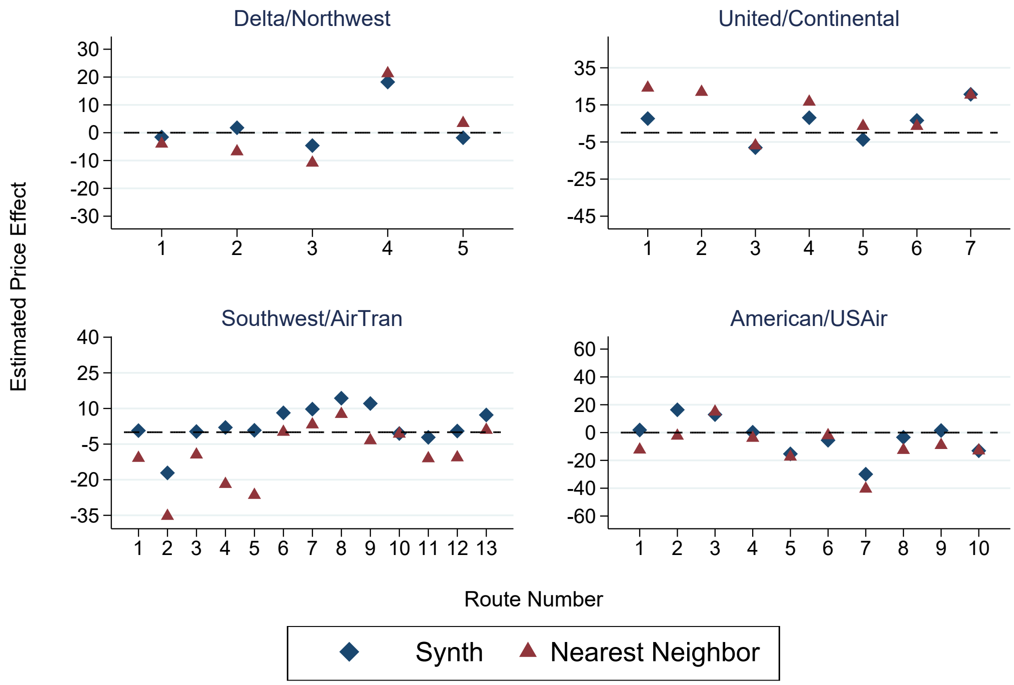

Notes: Each panel depicts the route-level estimated price effect for each of the four mergers. Blue diamonds represent the effect estimated using synthetic control. Red triangles represent the effect estimated using nearest neighbor. The vertical axis measures the coefficient estimates (multiplied by 100) and can be interpreted as the percentage change in price.

D Tables

Table 1: Route Summary Statistics

	DL/NW	UA/CO	WN/FL	AA/US
Overlap Routes	5	7	13	10
Legacy Carriers	2.00	2.43	0.08	2.40
	(0.00)	(0.53)	(0.28)	(0.52)
Legacy Share	0.84	0.96	0.03	0.99
	(0.15)	(0.10)	(0.12)	(0.03)
Pre-Merger HHI	4323	5052	5590	4865
	(885)	(1361)	(1068)	(715)
Δ HHI	3371	3021	3775	3693
	(1395)	(928)	(1169)	(1228)
Control - No Merging Firms	678	609	207	387
Legacy Carriers	0.70	0.60	1.39	0.39
	(0.69)	(0.65)	(0.81)	(0.53)
Legacy Share	0.46	0.36	0.71	0.22
	(0.46)	(0.42)	(0.71)	(0.34)
Pre-Merger HHI	8572	8130	7073	8277
	(2075)	(2223)	(2372)	(2236)
Control - One Merging Firm	176	153	250	161
Legacy Carriers	1.31	1.51	0.57	1.42
	(0.51)	(0.64)	(0.65)	(0.54)
Legacy Share	0.82	0.81	0.21	0.78
	(0.24)	(0.28)	(0.26)	(0.24)
Pre-Merger HHI	7117	6292	7296	6458
	(2331)	(2472)	(2494)	(2405)

Notes: Route-level summary statistics are provided for overlap routes, control routes with no merging firms present pre-merger, and control routes with one merging firm present pre-merger. Means and standard deviations are reported. Legacy carriers and legacy share present the number and quantity share of legacy airlines, respectively, on a route in the year prior to the merger. Pre-merger HHI calculates the Hirschman-Herfindahl Index in the year prior to the merger using quantity shares. Δ HHI uses the pre-merger shares to calculate the change in HHI that arises from the merger.

Table 2: Price and Passenger Summary Statistics

	Delta/Northwest			
	Pre-Merger	Post-Merger	Difference	Routes
Price - Overlap	179 (47)	174 (43)	-5.6	5
Price - Control 0	172 (52)	171 (49)	-0.9	678
Price - Control 1	205 (59)	188 (47)	-17.2***	176
Passengers - Overlap	5765 (1849)	5093 (1701)	-672**	5
Passengers - Control 0	3557 (93)	3094 (81)	-463***	678
Passengers - Control 1	3582 (190)	3370 (167)	-212***	176
	United/Continental			
	Pre-Merger	Post-Merger	Difference	Routes
Price - Overlap	185 (40)	214 (34)	28.3***	7
Price - Control 0	159 (43)	172 (44)	13.5***	609
Price - Control 1	185 (48)	202 (48)	17.0***	153
Passengers - Overlap	8840 (2092)	7724 (1540)	-1116	7
Passengers - Control 0	5118 (158)	5096 (156)	-22	609
Passengers - Control 1	5778 (323)	5700 (346)	-78	153
	Southwest/AirTran			
	Pre-Merger	Post-Merger	Difference	Routes
Price - Overlap	127 (8)	149 (13)	21.9***	13
Price - Control 0	180 (45)	193 (49)	12.6***	207
Price - Control 1	144 (30)	157 (32)	13.0***	250
Passengers - Overlap	10211 (1712)	9574 (5799)	-637*	13
Passengers - Control 0	7165 (311)	7560 (354)	395***	207
Passengers - Control 1	7281 (293)	7188 (293)	-93	250
	American/USAir			
	Pre-Merger	Post-Merger	Difference	Routes
Price - Overlap	205 (27)	192 (28)	-13.0*	10
Price - Control 0	170 (46)	173 (44)	3.1***	387
Price - Control 1	192 (44)	197 (44)	5.4***	161
Passengers - Overlap	7460 (1328)	9297 (1554)	1837***	10
Passengers - Control 0	5287 (182)	5472 (188)	185***	387
Passengers - Control 1	5895 (304)	6288 (332)	393***	161

Notes: This table summarizes pre and post-merger prices and passengers on overlap routes and both categories of control routes. Control group 1 are routes where neither merging firm was present per-merger and control group 2 are routes where exactly one merging firm was present pre-merger. We use the 8 pre and post merger quarters to calulate the route-level mean, and then report the mean across routes. Standard deviations are in parenthesis. The "Difference" column reports a paired t-test between the pre and post merger statistics. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 3: Synthetic Control Routes Summary

	DL/NW	UA/CO	WN/FL	AA/US
Overlap Routes				
Total Routes	5	6	13	10
Pre-Merger Price	175	194	126	202
Post-Merger Price	169	221	147	189
Pre-Merger Passengers	4217	7420	8765	6186
Post-Merger Passengers	3565	6510	8110	7781
Legacy Carriers	2.00	2.33	0.08	2.40
Legacy Share	0.84	0.96	0.03	0.99
Pre-Merger HHI	4323	5307	5590	4865
Control - No Merging Firms				
Pre-Merger Price	175	194	126	202
Post-Merger Price	165	210	143	196
Pre-Merger Passengers	2361	3364	7426	3429
Post-Merger Passengers	2014	3475	7127	3866
Legacy Carriers	0.91	1.08	0.65	0.83
Legacy Share	0.63	0.61	0.43	0.44
Pre-Merger HHI	8476	7860	8570	7469
Control - One Merging Firm				
Pre-Merger Price	175	194	126	202
Post-Merger Price	160	211	145	205
Pre-Merger Passengers	3268	4794	7470	5150
Post-Merger Passengers	3217	4625	6995	5229
Legacy Carriers	1.17	1.66	0.73	1.49
Legacy Share	0.82	0.83	0.26	0.81
Pre-Merger HHI	7282	5988	6989	6605

Notes: Route-level summary statistics are provided for overlap routes, the synthetic control routes constructed from routes with no merging firms present pre-merger, and the synthetic control routes constructed from routes with one merging firm present pre-merger. Prices and passengers are matched using logs, but are transformed to levels for ease of interpretation and comparison to other tables. Legacy carriers and legacy share present the number and quantity share of legacy airlines, respectively, on a route in the year prior to the merger. Pre-merger HHI calculates the HHI in the year prior to the merger using quantity shares. Pre-merger (post-merger) prices and passengers are the quarterly average across routes for the 8 pre-merger (post-merger) quarters.

Table 4: Nearest Neighbor Routes Summary

	DL/NW	UA/CO	WN/FL	AA/US
Overlap Routes				
Total Routes	5	7	13	10
Pre-Merger Price	175	180	124	202
Post-Merger Price	169	210	147	189
Pre-Merger Passengers	4217	6995	8765	6186
Post-Merger Passengers	3565	6414	8110	7781
Legacy Carriers	2.00	2.43	0.08	2.40
Legacy Share	0.84	0.96	0.03	0.99
Pre-Merger HHI	4323	5052	5590	4865
Control - No Merging Firms				
Pre-Merger Price	179	178	155	191
Post-Merger Price	161	195	168	193
Pre-Merger Passengers	1961	4384	6321	3788
Post-Merger Passengers	2132	4382	6814	3935
Legacy Carriers	1.04	0.81	1.29	0.41
Legacy Share	0.75	0.45	0.72	0.28
Pre-Merger HHI	8312	7699	7455	8698
Control - One Merging Firm				
Pre-Merger Price	174	182	139	199
Post-Merger Price	154	204	154	204
Pre-Merger Passengers	3529	5279	6672	4483
Post-Merger Passengers	3236	4953	6339	4650
Legacy Carriers	1.13	1.54	0.65	1.32
Legacy Share	0.66	0.88	0.29	0.81
Pre-Merger HHI	6125	6825	6992	6901

Notes: Route-level summary statistics are provided for overlap routes, nearest neighbor routes with no merging firms present pre-merger, and nearest neighbor routes constructed from routes with one merging firm present pre-merger. Prices and passengers are matched using logs, but are transformed to levels for ease of interpretation and comparison to other tables. Legacy carriers and legacy share present the number and quantity share of legacy airlines, respectively, on a route in the year prior to the merger. Pre-merger HHI calculates the HHI in the year prior to the merger using quantity shares. Pre-merger (post-merger) prices and passengers are the quarterly average across routes for the 8 pre-merger (post-merger) quarters.

Table 5: RMSE of Synthetic Control Models

	All Data	Yearly Avg	2-Quarter Avg	Staggered Quarters	No Controls
WNFL	0.160	0.105	0.089	0.059	0.160
AAUS	0.132	0.126	0.173	0.154	0.129
DLNW	0.052	0.052	0.061	0.089	0.129
UACO	0.098	0.094		0.062	0.098
Method Average	0.111	0.094	0.108	0.091	0.129

Notes: Root mean square error is computed over the 8 backwards forecast pre-treatment periods (periods 16-8 before treatment) from models fit on the 8 quarters pre-treatment. The "All Data" model is fit using data for every quarter pre-treatment. The "Yearly Avg" model uses the yearly average of prices in the 2 years pre-treatment. Similarly, "2-Quarter Avg" uses average of two quarter increments for the 8 quarters before treatment. Lastly, "Staggered Quarters" uses data from quarters 8-, 6-, 3-, and 1- quarter pre-merger. The "2-Quarter Avg" for UACO is omitted as the model did not converge. The method average is the average RMSE for the method across mergers.

Table 6: Summary of Price Effect Estimates

	DL/NW	UA/CO	WN/FL	AA/US
OLS	-0.035	0.069	0.093	-0.089
	(99% - First Year)	(95% - First Year)	(99%)	(95%)
Synth Control	0.024	0.052	0.028	-0.034
	(None)	(99% - one quarter)	(No Effect)	(95% - two quarters)
Synth Control (add. controls)	-0.029	0.031	0.078	$0.013^{\dagger}$
	(95% - two quarters)	(90% - one quarter)	(90% - two quarters)	(95% - two quarters)
Nearest Neighbor	0.006	0.119	-0.091	-0.098
	(None)	(99%)	(99%)	(95%)
Nearest Neighbor (add. controls)	-0.011	0.072	$-0.059^{\dagger\dagger}$	-0.032
	(None)	(95%)	(99%)	(95% - two quarters)

Notes: This table summarizes the price effect estimates by merger and DID methodology. For synthetic control and nearest neighbor, (add. controls) is the specification that adds competition controls. Notes on statistical significance are in parenthesis below the estimate. All estimates are relative to control group 0, where no merging firm is present pre-merger. For OLS the estimate is for the average effect across the two post-merger years. If there is statistical significance in one of the two post-merger years, but when pooling both years, that is reported. For synthetic control and nearest neighbor, price effects are separately estimated for 8 post-merger quarters. The table reports the average estimate across all 8 quarters. If at least 3 quarters are statistically significant, then the highest level of significance is reported. If fewer than 3 quarters are significant than the level and number of significant quarters are reported. † = Although the average effect is positive, the two quarters that are statistically significant are negative. †† = The average effect is negative, but one of the four statistically significant quarters is estimated to be positive.

D.1 Delta-Northwest Tables

Table 7: DL-NW Merger: Price Regression

	(1)	(2)	(3)	(4)
Merger X Post	-0.035		0.043	
	(0.03)		(0.04)	
Merger X Post1		-0.049***		0.031*
		(0.01)		(0.02)
Merger X Post2		-0.021		0.054
		(0.06)		(0.06)
Constant	5.099***	5.099***	5.234***	5.234***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	10928	10928	2896	2896
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No DL/NW	No DL/NW	DL/NW: No Overlap	DL/NW: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 8: DL-NW Merger: Price Synthetic Control

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.036	0.420	0.081	0.122	-0.002	0.402	-0.003	0.418
T+2	-0.050	0.102	-0.026	0.280	-0.040	0.233	-0.041	0.257
T+3	-0.035	0.157	-0.094	0.041	0.055	0.302	0.055	0.320
T+4	0.048	0.205	-0.014	0.359	0.061	0.270	0.061	0.270
T+5	0.080	0.218	0.036	0.440	0.060	0.288	0.060	0.279
T+6	0.056	0.381	-0.014	0.285	0.087	0.190	0.087	0.189
T+7	0.056	0.386	-0.070	0.103	0.110	0.139	0.110	0.137
T+8	0.001	0.322	-0.133	0.014	0.081	0.178	0.080	0.181
Treated Routes	5		4		5		5	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route. The "Prices + Controls" specification uses prices for all pre-merger quarters in addition to the competition variables detailed in the body of the paper to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 9: DL-NW Merger: Price Nearest Neighbor

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.048	0.219	0.024	0.608	0.062	0.304	0.071	0.229
T+2	-0.015	0.602	-0.047	0.261	0.018	0.716	-0.007	0.909
T+3	-0.020	0.770	-0.026	0.698	0.069	0.004	0.063	0.143
T+4	0.037	0.507	0.034	0.645	0.070	0.000	0.096	0.015
T+5	0.019	0.747	0.013	0.834	0.081	0.000	0.062	0.106
T+6	-0.016	0.859	0.011	0.897	0.114	0.012	0.043	0.364
T+7	0.007	0.956	0.004	0.971	0.144	0.021	0.096	0.121
T+8	-0.009	0.924	-0.099	0.256	0.099	0.004	0.121	0.050
Treated Routes	5		5		5		5	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to select the 4 nearest neighbors for each overlap route. The "Prices + Controls" specification uses prices for all pre-merger quarters in addition to the competition variables detailed in the body of the paper to select the 4 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

D.2 United-Continental Tables

Table 10: UA-CO Merger: Price Regression

	(1)	(2)	(3)	(4)
Merger X Post	0.069*		0.062	
	(0.04)		(0.04)	
Merger X Post1		0.066**		0.068**
		(0.03)		(0.03)
Merger X Post2		0.072		0.055
		(0.05)		(0.06)
Constant	5.074***	5.074***	5.233***	5.233***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	9856	9856	2560	2560
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No UA/CO	No UA/CO	UA/CO: No Overlap	UA/CO: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The data include observations for two years before and after the merger. The "No UA/CO" control group includes only routes with no UA or CO direct flights. The "UA/CO: No Overlap" control group includes only routes with either UA or CO direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 11: UA-CO Merger: Price Synthetic Control

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.092	0.007	0.048	0.093	0.048	0.105	0.031	0.198
T+2	0.057	0.110	0.019	0.344	0.016	0.302	-0.001	0.471
T+3	0.042	0.134	0.026	0.250	0.073	0.027	0.103	0.022
T+4	0.051	0.138	0.030	0.304	0.079	0.064	0.147	0.016
T+5	0.019	0.367	0.022	0.333	0.019	0.440	0.079	0.196
T+6	0.050	0.249	0.016	0.416	0.112	0.043	0.136	0.061
T+7	0.032	0.243	0.024	0.270	0.075	0.118	0.110	0.065
T+8	0.074	0.115	0.064	0.108	0.084	0.099	0.131	0.047
Treated Routes	6		6		7		7	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route. The "Prices + Controls" specification uses prices for a subset of pre-merger quarters and the competition variables detailed in the body of the paper to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 12: UA-CO Merger: Price Nearest Neighbor

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.119	0.000	0.073	0.032	0.078	0.010	0.091	0.007
T+2	0.116	0.026	0.064	0.244	0.057	0.148	0.055	0.231
T+3	0.114	0.042	0.083	0.078	0.046	0.197	0.023	0.475
T+4	0.139	0.006	0.085	0.076	0.026	0.276	0.013	0.695
T+5	0.110	0.019	0.057	0.233	-0.006	0.833	-0.019	0.547
T+6	0.140	0.068	0.103	0.098	0.007	0.857	0.058	0.373
T+7	0.086	0.159	0.041	0.454	-0.041	0.401	-0.018	0.786
T+8	0.129	0.013	0.067	0.242	0.002	0.966	0.013	0.851
Treated Routes	7		7		7		7	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to select the 5 nearest neighbors for each overlap route. The "Prices + Controls" specification uses prices for all pre-merger quarters in addition to the competition variables detailed in the body of the paper to select the 5 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

D.3 Southwest-Airtran Tables

Table 13: WN-FL Merger: Price Regression

	(1)	(2)	(3)	(4)
Merger X Post	0.093***		0.075***	
	(0.02)		(0.02)	
Merger X Post1		0.097***		0.082***
		(0.02)		(0.01)
Merger X Post2		0.090***		0.067**
		(0.03)		(0.03)
Constant	5.170***	5.170***	4.984***	4.984***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	3520	3520	4208	4208
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No WN/FL	No WN/FL	WN/FL: No Overlap	WN/FL: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The data include observations for two years before and after the merger. The "No WN/FL" control group includes only routes with no WN or FL direct flights. The "WN/FL: No Overlap" control group includes only routes with either WN or FL direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 14: WN-FL Merger: Price Synthetic Control

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.004	0.244	0.017	0.354	-0.012	0.147	-0.029	0.083
T+2	0.014	0.364	0.077	0.182	0.017	0.028	0.031	0.010
T+3	0.042	0.445	0.129	0.059	0.016	0.100	0.055	0.003
T+4	0.053	0.288	0.098	0.078	0.023	0.049	0.074	0.001
T+5	0.041	0.387	0.063	0.218	0.028	0.081	0.026	0.102
T+6	0.029	0.499	0.086	0.141	0.030	0.045	0.023	0.050
T+7	0.035	0.484	0.103	0.114	0.038	0.009	0.050	0.002
T+8	0.012	0.338	0.047	0.498	0.012	0.239	-0.010	0.390
Treated Routes	13		13		13		11	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route. The "Prices + Controls" specification uses price for a subset of pre-merger quarters and the competition variables detailed in the body of the paper to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 15: WN-FL Merger: Price Nearest Neighbor

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.190	0.000	-0.057	0.303	-0.009	0.687	0.062	0.041
T+2	-0.143	0.000	-0.025	0.649	0.027	0.220	-0.014	0.667
T+3	-0.165	0.000	0.181	0.004	0.040	0.096	-0.009	0.770
T+4	-0.092	0.011	0.012	0.852	0.053	0.057	0.090	0.004
T+5	-0.027	0.640	-0.080	0.265	0.047	0.179	0.126	0.001
T+6	-0.008	0.888	-0.193	0.004	0.065	0.114	0.058	0.222
T+7	-0.072	0.221	-0.175	0.006	0.055	0.130	0.083	0.048
T+8	-0.028	0.655	-0.133	0.054	0.029	0.498	0.107	0.021
Treated Routes	13		13		13		13	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to select the 5 nearest neighbors for each overlap route. The "Prices + Controls" specification uses prices for all pre-merger quarters in addition to the competition variables detailed in the body of the paper to select the 5 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

D.4 American Airlines-USAir Tables

Table 16: AA-US Merger: Price Regression

	(1)	(2)	(3)	(4)
Merger X Post	-0.089**		-0.097***	
	(0.04)		(0.04)	
Merger X Post1		-0.069**		-0.092***
		(0.03)		(0.03)
Merger X Post2		-0.109**		-0.101**
		(0.05)		(0.05)
Constant	5.115***	5.115***	5.246***	5.246***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	6352	6352	2736	2736
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No AA/US	No AA/US	AA/US: No Overlap	AA/US: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The data include observations for two years before and after the merger. The "No AA/US" control group includes only routes with no AA or US direct flights. The "AA/US: No Overlap" control group includes only routes with either AA or US direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 17: AA-US Merger: Price Synthetic Control

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.060	0.016	-0.069	0.005	-0.072	0.013	-0.085	0.032
T+2	-0.030	0.120	0.003	0.390	-0.073	0.012	-0.087	0.020
T+3	-0.017	0.255	0.051	0.220	-0.084	0.009	-0.087	0.025
T+4	-0.031	0.123	0.013	0.376	-0.071	0.058	-0.066	0.058
T+5	-0.022	0.219	0.014	0.363	-0.071	0.130	-0.057	0.229
T+6	0.020	0.424	0.080	0.126	-0.062	0.262	-0.055	0.422
T+7	-0.044	0.102	0.065	0.290	-0.097	0.159	-0.091	0.242
T+8	-0.091	0.011	-0.056	0.050	-0.109	0.086	-0.115	0.105
Treated Routes	10		9		10		10	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route. The "Prices + Controls" specification uses annual average pre-merger prices and the competition variables detailed in the body of the paper to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 18: AA-US Merger: Price Nearest Neighbor

Control Group:	No Merging Airlines				One Merging Airline			
Match On:	Prices		Prices + Controls		Prices		Prices + Controls	
	Effect	P-Value	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.102	0.011	-0.108	0.032	-0.085	0.055	-0.081	0.077
T+2	-0.068	0.068	-0.073	0.046	-0.068	0.054	-0.085	0.045
T+3	-0.066	0.076	0.001	0.983	-0.086	0.019	-0.104	0.029
T+4	-0.077	0.092	0.008	0.840	-0.084	0.030	-0.091	0.078
T+5	-0.104	0.050	-0.043	0.530	-0.082	0.089	-0.089	0.245
T+6	-0.060	0.267	0.026	0.726	-0.065	0.172	-0.072	0.380
T+7	-0.141	0.086	-0.020	0.859	-0.107	0.192	-0.141	0.252
T+8	-0.168	0.019	-0.045	0.642	-0.126	0.079	-0.132	0.238
Treated Routes	10		10		10		10	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Prices" specification uses the 8 pre-merger quarters to select the 5 nearest neighbors for each overlap route. The "Prices + Controls" specification uses prices for all pre-merger quarters in addition to the competition variables detailed in the body of the paper to select the 5 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

D.5 Multi-Market Contact

Table 19: Multi-Market Contact Summary Statistics

	Average Multi-Market Contact			
Merger	Pre-Merger	Post-Merger	Difference	Routes
Delta/Northwest	30.0 (18.7)	35.8 (15.1)	5.8***	82
United/Continental	38.3 (25.8)	50.3 (26.8)	12.0***	97
Southwest/AirTran	38.0 (15.6)	51.7 (20.5)	13.7***	130
American/USAir	44.3 (22.3)	62.5 (20.9)	18.2***	99

Notes: This table summarizes the average pre and post-merger multi-market contact on routes with only one merging firm present pre-merger (control group 2). Only routes with at least two airlines present for all 8 pre and post-merger quarters are included. The mean values across routes is reported and the standard deviations are in parenthesis. The "Difference" column reports a paired t-test between the pre and post merger means. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 20: Regression of Log Price on Change in Multi-Market Contact

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMMC x Post	0.521*	-0.086	0.155*	0.170**	0.111**	0.165***	0.163*	0.149**
	(0.29)	(0.11)	(0.09)	(0.07)	(0.05)	(0.04)	(0.09)	(0.07)
Constant	508***	509***	517***	510***	499***	508***	522***	517***
	(0.72)	(0.02)	(0.58)	(0.25)	(0.38)	(0.20)	(0.85)	(0.31)
Observations	1312	3424	1552	4368	2080	3808	1584	3296
Merger	DL/NW	DL/NW	UA/CO	UA/CO	SW/AT	SW/AT	AA/US	AA/US
Routes	One-Firm Control	All Control	One-Firm Control	All Control	One-Firm Control	All Control	One-Firm Control	All Control
Fixed Effect	Route/Time	Route/Time	Route/Time	Route/Time	Route/Time	Route/Time	Route/Time	Route/Time

Notes: Data is at the route-quarter level. The dependent variable is log-price (multiplied by 100 for table readability). Only routes with at least two airlines for all 8 pre- and post-merger quarters are included. ΔMMC x Post is the change in average route-level multi-market contact from the year before to the year after the merger, interacted with a post-merger indicator. "One-Firm Control" denotes routes with only one merging firm present pre-merger. "All Control" denotes routes with no merging firm present pre-merger. Standard errors are clustered at the route-quarter level.

D.6 Route-Level Effects

Table 21: Determinants of Route-Level Effects

	Synth	NN
HHI	3.09**	0.75
	[2.18, 6.19]	[-1.15, 2.86]
Legacy Carriers	0.00	6.27**
	[-0.54, 1.93]	[4.95, 12.54]
Circuity	-4.19**	-4.31**
	[-8.26, -3.35]	[-8.61, -2.38]
Passengers	-1.67**	-3.67**
	[-3.34, -0.45]	[-7.34, -0.67]
Unemployment	0.00	1.88
	[-1.82, 1.17]	[-0.29, 4.94]
Income	0.00	-2.96**
	[-0.39, 2.02]	[-6.46, -1.36]
Population	0.00	0.10
	[-0.92, 1.66]	[-2.40, 2.69]
Miles	-2.63**	0.00
	[-5.27, -1.98]	[-2.76, 2.53]
DHHI	-3.06**	-1.80
	[-6.12, -2.12]	[-3.75, 1.25]
Carriers	0.00	0.00
	[0.00, 0.00]	[-4.47, 0.00]
Observations	34	35

Notes: Coefficient estimates from a leave-one-out cross-validated Lasso procedure are presented. The bracketed numbers are 95% confidence intervals from a 500 iteration bootstrap of the Lasso estimation. Leave-one-out cross-validation is used to select the optimal lambda tuning parameter. A '**' indicates that a parameter has 95% confidence interval that exclude 0.

E Appendix Figures and Tables

Figure 9: Prices Indices

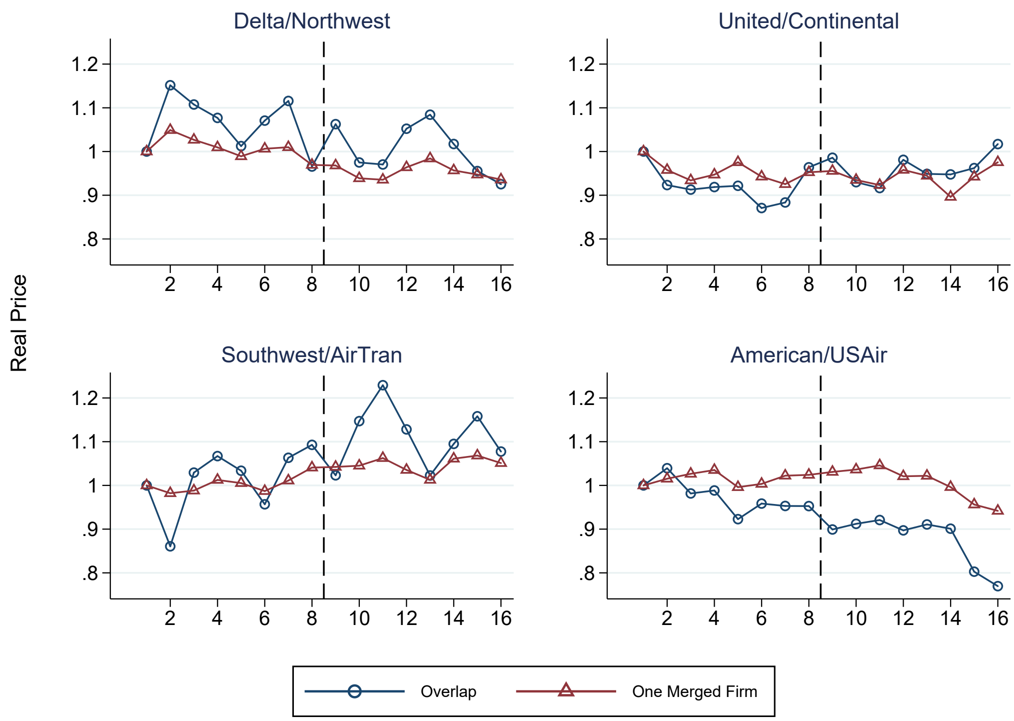

Notes: Each panel depicts a quarterly-average price index for 8 pre and post-merger quarters. The dashed line indicates the time at which the merger was consummated. Overlap indicates routes where both mergers competed prior to the merger. One Merged Firm indicates routes where exactly one firm competed pre-merger. The price index on overlap routes is created by first calculated the quarterly-average prices on overlaps and routes with no merged firms. Then, average price on overlaps is divided by the average price on no merged firm routes. Finally we normalize the index to be one in the earliest period. The index on one merged firm routes is calculated in the same way.

E.1 Unobserved Passenger Comparison

Table 22: RMSE of Unobserved Passenger Trends

Method:	OLS		Nearest Neighbor		Synthetic Control	
Control Group:	No Merging Airlines	One Merging Airline	No Merging Airlines	One Merging Airline	No Merging Airlines	One Merging Airline
UACO	3101	2466	2682	1766	4123	2707
AAUS	1877	1414	2417	1724	2780	1075
DLNW	1478	1536	2298	765	1888	1052
WNFL	2819	2700	2522	2125	1409	1411
Average:	2319	2029	2480	1595	2550	1561
Method Average:	2174		2037		2056	

Notes: The weights of the control group, fit using pre-merger log-prices, are used to calculate the number of log-passengers for every quarter pre-merger in the control group. The root mean squared error between the control group log-passenger trend and the treated group log-passenger trend is then calculated. We present the RMSE in levels for each method and control group in every merger.

Table 23: RMSE of Unobserved Passenger Trends: Additional Controls

Method:	Nearest Neighbor		NN + Controls		Synthetic Control		Optimal SC	
Control Group:	No Merging Airlines	One Merging Airline	No Merging Airlines	One Merging Airline	No Merging Airlines	One Merging Airline	No Merging Airlines	One Merging Airline
UACO	2682	1766	627	493	4123	2707	3923	3033
AAUS	2417	1724	1572	624	2780	1075	2733	1255
DLNW	2298	765	588	898	1888	1052	1535	1038
WNFL	2522	2125	1478	2772	1409	1411	836	3282
Average:	2480	1595	1066	1197	2550	1561	2257	2152
Method Average:	2037		1132		2056		2204	

Notes: The weights of the control group, fit using pre-merger log-prices, are used to calculate the number of log-passengers for every quarter pre-merger in the control group. The root mean squared error between the control group log-passenger trend and the treated group log-passenger trend is then calculated. We present the RMSE in levels for each method and control group in every merger.

F Web Appendix - Passenger Results

F.1 Delta/Northwest

Table 24: DL-NW Merger: Passengers Regression

	(1)	(2)	(3)	(4)
Merger X Post	-0.027		-0.136**	
	(0.05)		(0.06)	
Merger X Post1		0.041*		-0.060**
		(0.02)		(0.03)
Merger X Post2		-0.094		-0.211*
		(0.11)		(0.11)
Constant	7.866***	7.866***	7.910***	7.910***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	10928	10928	2896	2896
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No DL/NW	No DL/NW	DL/NW: No Overlap	DL/NW: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the total non-stop passengers. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level.

Table 25: DL-NW Merger: Passengers Synthetic Control

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	0.076	0.149	-0.017	0.490
T+2	0.041	0.199	-0.048	0.461
T+3	0.018	0.259	-0.131	0.242
T+4	0.008	0.289	-0.227	0.106
T+5	-0.112	0.365	-0.268	0.080
T+6	-0.129	0.303	-0.315	0.043
T+7	-0.196	0.156	-0.369	0.038
T+8	0.016	0.284	-0.211	0.190
Treated Routes	5		5	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 26: DL-NW Merger: Passenger Nearest Neighbor

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	0.162	0.014	-0.087	0.381
T+2	0.210	0.000	-0.026	0.580
T+3	0.166	0.023	-0.061	0.389
T+4	0.107	0.068	-0.171	0.023
T+5	0.044	0.096	-0.243	0.057
T+6	0.036	0.725	-0.229	0.113
T+7	-0.036	0.825	-0.342	0.147
T+8	0.129	0.079	-0.139	0.167
Treated Routes	5		5	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to select the 4 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

F.2 United/Continental

Table 27: UA-CO Merger: Passengers Regression

	(1)	(2)	(3)	(4)
Merger X Post	-0.083		-0.037	
	(0.06)		(0.07)	
Merger X Post1		-0.063		-0.042
		(0.04)		(0.05)
Merger X Post2		-0.103		-0.032
		(0.09)		(0.09)
Constant	8.294***	8.294***	8.428***	8.428***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	9856	9856	2560	2560
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No UA/CO	No UA/CO	UA/CO: No Overlap	UA/CO: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the total non-stop passengers. The data include observations for two years before and after the merger. The "No UA/CO" control group includes only routes with no UA or CO direct flights. The "UA/CO: No Overlap" control group includes only routes with either UA or CO direct flights, but not both. Standard errors are clustered at the route level.

Table 28: UA-CO Merger: Passengers Synthetic Control

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	0.005	0.327	-0.028	0.253
T+2	-0.087	0.241	-0.150	0.048
T+3	0.006	0.249	-0.047	0.276
T+4	0.037	0.208	-0.043	0.291
T+5	0.009	0.373	-0.052	0.286
T+6	-0.116	0.237	-0.127	0.169
T+7	-0.056	0.415	-0.117	0.154
T+8	0.018	0.389	-0.100	0.220
Treated Routes	7		6	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 29: UA-CO Merger: Passenger Nearest Neighbor

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	-0.121	0.008	-0.050	0.445
T+2	-0.268	0.000	-0.095	0.397
T+3	-0.092	0.171	-0.099	0.193
T+4	-0.091	0.272	-0.164	0.107
T+5	-0.155	0.030	-0.097	0.393
T+6	-0.174	0.029	-0.105	0.419
T+7	-0.105	0.193	-0.144	0.345
T+8	-0.010	0.932	-0.137	0.456
Treated Routes	7		7	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to select the 5 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

F.3 Southwest/AirTrain

Table 30: WN-FL Merger: Passengers Regression

	(1)	(2)	(3)	(4)
Merger X Post	-0.105***		-0.061**	
	(0.03)		(0.03)	
Merger X Post1		-0.110***		-0.084***
		(0.02)		(0.02)
Merger X Post2		-0.101**		-0.037
		(0.05)		(0.05)
Constant	8.739***	8.739***	8.730***	8.730***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	3520	3520	4208	4208
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No WN/FL	No WN/FL	WN/FL: No Overlap	WN/FL: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the total non-stop passengers. The data include observations for two years before and after the merger. The "No WN/FL" control group includes only routes with no WN or FL direct flights. The "WN/FL: No Overlap" control group includes only routes with either WN or FL direct flights, but not both. Standard errors are clustered at the route level.

Table 31: WN-FL Merger: Passengers Synthetic Control

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	-0.188	0.000	-0.119	0.005
T+2	-0.103	0.019	-0.109	0.001
T+3	-0.030	0.454	-0.045	0.212
T+4	-0.124	0.060	-0.023	0.421
T+5	-0.222	0.006	-0.119	0.006
T+6	-0.059	0.174	0.007	0.324
T+7	-0.010	0.494	0.041	0.120
T+8	-0.079	0.123	-0.013	0.426
Treated Routes	13		11	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 32: WN-FL Merger: Passenger Nearest Neighbor

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	-0.118	0.020	-0.074	0.146
T+2	-0.143	0.009	-0.156	0.004
T+3	-0.091	0.089	-0.143	0.002
T+4	-0.111	0.049	-0.138	0.006
T+5	-0.101	0.149	-0.189	0.000
T+6	0.004	0.955	-0.120	0.077
T+7	0.019	0.792	-0.086	0.259
T+8	-0.017	0.827	-0.106	0.144
Treated Routes	13		13	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to select the 5 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

F.4 American/USAir

Table 33: AA-US Merger: Passengers Regression

	(1)	(2)	(3)	(4)
Merger X Post	0.203***		0.176***	
	(0.06)		(0.06)	
Merger X Post1		0.175***		0.170***
		(0.05)		(0.05)
Merger X Post2		0.231***		0.183***
		(0.07)		(0.07)
Constant	8.398***	8.398***	8.518***	8.518***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	6352	6352	2736	2736
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No AA/US	No AA/US	AA/US: No Overlap	AA/US: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the total non-stop passengers. The data include observations for two years before and after the merger. The "No AA/US" control group includes only routes with no AA or US direct flights. The "AA/US: No Overlap" control group includes only routes with either AA or US direct flights, but not both. Standard errors are clustered at the route level.

Table 34: AA-US Merger: Passengers Synthetic Control

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	0.095	0.023	0.103	0.013
T+2	0.101	0.014	0.096	0.017
T+3	0.095	0.033	0.083	0.048
T+4	0.104	0.027	0.126	0.013
T+5	0.104	0.040	0.150	0.024
T+6	0.111	0.028	0.114	0.111
T+7	0.166	0.013	0.111	0.152
T+8	0.150	0.012	0.106	0.141
Treated Routes	10		10	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to construct a synthetic control route for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 35: AA-US Merger: Passenger Nearest Neighbor

Control Group:	No Merging Airlines		One Merging Airline	
Match On:	Passengers		Passengers	
	Effect	P-Value	Effect	P-Value
T+1	0.117	0.063	0.113	0.073
T+2	0.100	0.127	0.133	0.031
T+3	0.071	0.317	0.131	0.046
T+4	0.088	0.154	0.143	0.012
T+5	0.071	0.326	0.127	0.076
T+6	0.087	0.341	0.121	0.157
T+7	0.141	0.259	0.110	0.267
T+8	0.140	0.181	0.061	0.403
Treated Routes	10		10	

Notes: Data is at the route-quarter level. The dependent variable is the log of total passengers. The log of total passengers in the 8 pre-merger quarters are used to select the 5 nearest neighbors for each overlap route. "T+X" is the estimate for the Xth quarter after the merger occurred.

G Web Appendix - No Monopoly and No Entry/Exit Results

G.1 Delta/Northwest

Table 36: DL-NW Merger: No Monopoly Routes

	(1)	(2)	(3)	(4)
Merger X Post	-0.024		0.014	
	(0.04)		(0.04)	
Merger X Post1		-0.033**		0.005
		(0.02)		(0.02)
Merger X Post2		-0.015		0.024
		(0.06)		(0.06)
Constant	5.111***	5.111***	5.108***	5.108***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	2816	2816	1520	1520
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No DL/NW	No DL/NW	DL/NW: No Overlap	DL/NW: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes with only one airline present for all 8 pre-merger quarters are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 37: DL-NW Merger: No Entry/Exit Routes

	(1)	(2)	(3)	(4)
Merger X Post	-0.035		0.043	
	(0.03)		(0.04)	
Merger X Post1		-0.049***		0.031*
		(0.01)		(0.02)
Merger X Post2		-0.021		0.054
		(0.06)		(0.06)
Constant	5.099***	5.099***	5.234***	5.234***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	10928	10928	2896	2896
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No DL/NW	No DL/NW	DL/NW: No Overlap	DL/NW: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes that that experienced post-merger entry or exit are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 38: DL-NW Merger: Price Synthetic Control Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.036	0.449	0.053	0.273	0.036	0.447
T+2	-0.050	0.107	-0.019	0.228	-0.050	0.105
T+3	-0.035	0.172	0.029	0.282	-0.035	0.168
T+4	0.048	0.224	0.078	0.080	0.048	0.225
T+5	0.080	0.219	0.080	0.147	0.080	0.230
T+6	0.056	0.368	0.025	0.467	0.056	0.370
T+7	0.056	0.378	0.014	0.480	0.056	0.374
T+8	0.001	0.332	-0.033	0.232	0.001	0.324
Treated Routes	5		5		5	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route and matches on only pre-merger prices. Control routes only include routes where no merging firm was present pre-merger. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 39: DL-NW Merger: Price Nearest Neighbor Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.048	0.219	0.052	0.053	0.048	0.219
T+2	-0.015	0.602	-0.031	0.261	-0.015	0.602
T+3	-0.020	0.770	-0.042	0.499	-0.020	0.770
T+4	0.037	0.507	0.049	0.471	0.037	0.507
T+5	0.019	0.747	0.057	0.318	0.019	0.747
T+6	-0.016	0.859	0.021	0.818	-0.016	0.859
T+7	0.007	0.956	-0.001	0.994	0.007	0.956
T+8	-0.009	0.924	-0.073	0.451	-0.009	0.924
Treated Routes	5		5		5	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the control group with no merging firms present to select the five nearest neighbors. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

G.2 United/Continental

Table 40: UA-CO Merger: No Monopoly Routes

	(1)	(2)	(3)	(4)
Merger X Post	0.066*		0.050	
	(0.04)		(0.04)	
Merger X Post1		0.070**		0.059*
		(0.03)		(0.03)
Merger X Post2		0.061		0.042
		(0.05)		(0.06)
Constant	5.055***	5.055***	5.184***	5.184***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	3408	3408	1728	1728
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No UA/CO	No UA/CO	UA/CO: No Overlap	UA/CO: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes with only one airline present for all 8 pre-merger quarters are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 41: UA-CO Merger: No Entry/Exit Routes

	(1)	(2)	(3)	(4)
Merger X Post	0.099***		0.092***	
	(0.03)		(0.03)	
Merger X Post1		0.084***		0.086***
		(0.03)		(0.03)
Merger X Post2		0.114***		0.097**
		(0.04)		(0.04)
Constant	5.073***	5.073***	5.232***	5.232***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	9840	9840	2544	2544
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No UA/CO	No UA/CO	UA/CO: No Overlap	UA/CO: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes that that experienced post-merger entry or exit are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 42: UA-CO Merger: Price Synthetic Control Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.092	0.006	0.070	0.001	0.089	0.014
T+2	0.057	0.106	0.033	0.124	0.066	0.086
T+3	0.042	0.160	0.034	0.148	0.059	0.098
T+4	0.051	0.156	0.059	0.041	0.081	0.070
T+5	0.019	0.382	0.025	0.236	0.046	0.219
T+6	0.050	0.243	0.056	0.178	0.098	0.088
T+7	0.032	0.236	0.025	0.243	0.072	0.099
T+8	0.074	0.120	0.049	0.202	0.119	0.048
Treated Routes	6		7		5	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route and matches on only pre-merger prices. Control routes only include routes where no merging firm was present pre-merger. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 43: UA-CO Merger: Price Nearest Neighbor Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	0.119	0.000	0.098	0.007	0.118	0.001
T+2	0.116	0.026	0.086	0.141	0.135	0.025
T+3	0.114	0.042	0.088	0.077	0.150	0.015
T+4	0.139	0.006	0.093	0.011	0.188	0.000
T+5	0.110	0.019	0.058	0.154	0.153	0.000
T+6	0.140	0.068	0.116	0.095	0.198	0.007
T+7	0.086	0.159	0.051	0.385	0.143	0.012
T+8	0.129	0.013	0.088	0.047	0.183	0.000
Treated Routes	7		7		6	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the control group with no merging firms present to select the five nearest neighbors. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

G.3 Southwest/AirTran

Table 44: WN-FL Merger: No Monopoly Routes

	(1)	(2)	(3)	(4)
Merger X Post	0.081***		0.079***	
	(0.02)		(0.02)	
Merger X Post1		0.089***		0.085***
		(0.02)		(0.01)
Merger X Post2		0.072**		0.072**
		(0.03)		(0.03)
Constant	5.173***	5.173***	4.983***	4.983***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	2016	2016	2384	2384
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No WN/FL	No WN/FL	WN/FL: No Overlap	WN/FL: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes with only one airline present for all 8 pre-merger quarters are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 45: WN-FL Merger: No Entry/Exit Routes

	(1)	(2)	(3)	(4)
Merger X Post	0.104***		0.085***	
	(0.02)		(0.02)	
Merger X Post1		0.101***		0.086***
		(0.02)		(0.02)
Merger X Post2		0.106***		0.084***
		(0.02)		(0.02)
Constant	5.173***	5.173***	4.984***	4.984***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	3472	3472	4160	4160
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No WN/FL	No WN/FL	WN/FL: No Overlap	WN/FL: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes that that experienced post-merger entry or exit are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 46: WN-FL Merger: Price Synthetic Control Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.004	0.235	-0.004	0.447	0.004	0.319
T+2	0.014	0.369	0.032	0.050	0.012	0.375
T+3	0.042	0.455	0.097	0.002	0.038	0.491
T+4	0.053	0.296	0.103	0.002	0.044	0.396
T+5	0.041	0.408	0.130	0.000	0.053	0.326
T+6	0.029	0.481	0.156	0.000	0.047	0.364
T+7	0.035	0.499	0.117	0.001	0.048	0.392
T+8	0.012	0.335	0.073	0.019	0.025	0.451
Treated Routes	13		13		10	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route and matches on only pre-merger prices. Control routes only include routes where no merging firm was present pre-merger. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 47: WN-FL Merger: Price Nearest Neighbor Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.190	0.000	-0.097	0.035	-0.156	0.000
T+2	-0.143	0.000	-0.106	0.013	-0.114	0.004
T+3	-0.165	0.000	-0.029	0.511	-0.160	0.001
T+4	-0.092	0.011	-0.059	0.164	-0.073	0.039
T+5	-0.027	0.640	0.031	0.599	0.036	0.521
T+6	-0.008	0.888	0.049	0.425	0.023	0.681
T+7	-0.072	0.221	0.030	0.623	-0.045	0.393
T+8	-0.028	0.655	0.002	0.965	-0.005	0.929
Treated Routes	13		13		10	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the control group with no merging firms present to select the five nearest neighbors. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

G.4 American/USAir

Table 48: AA-US Merger: No Monopoly Routes

	(1)	(2)	(3)	(4)
Merger X Post	-0.092**		-0.100***	
	(0.04)		(0.04)	
Merger X Post1		-0.075**		-0.102***
		(0.04)		(0.04)
Merger X Post2		-0.108**		-0.099**
		(0.05)		(0.05)
Constant	5.135***	5.135***	5.238***	5.238***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	2000	2000	1792	1792
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No AA/US	No AA/US	AA/US: No Overlap	AA/US: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes with only one airline present for all 8 pre-merger quarters are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 49: AA-US Merger: No Entry/Exit Routes

	(1)	(2)	(3)	(4)
Merger X Post	-0.045		-0.053	
	(0.05)		(0.05)	
Merger X Post1		-0.066		-0.089*
		(0.05)		(0.05)
Merger X Post2		-0.024		-0.016
		(0.05)		(0.05)
Constant	5.113***	5.113***	5.244***	5.244***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	6288	6288	2672	2672
Route FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Control Group	No AA/US	No AA/US	AA/US: No Overlap	AA/US: No Overlap

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. Routes that that experienced post-merger entry or exit are dropped. The data include observations for two years before and after the merger. The "No DL/NW" control group includes only routes with no DL or NW direct flights. The "DL/NW: No Overlap" control group includes only routes with either DL or NW direct flights, but not both. Standard errors are clustered at the route level. ***, **, and * denote 1%, 5%, and 10% significance levels.

Table 50: AA-US Merger: Price Synthetic Control Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.060	0.015	-0.018	0.225	-0.033	0.121
T+2	-0.030	0.127	-0.015	0.284	-0.034	0.148
T+3	-0.017	0.265	0.006	0.482	-0.002	0.448
T+4	-0.031	0.143	0.005	0.470	-0.013	0.325
T+5	-0.022	0.231	0.015	0.415	0.042	0.312
T+6	0.020	0.411	0.033	0.342	0.078	0.104
T+7	-0.044	0.109	-0.054	0.111	0.091	0.122
T+8	-0.091	0.014	-0.084	0.058	0.012	0.391
Treated Routes	10		10		6	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the 8 pre-merger quarters to construct a synthetic control route for each overlap route and matches on only pre-merger prices. Control routes only include routes where no merging firm was present pre-merger. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.

Table 51: AA-US Merger: Price Nearest Neighbor Robustness

	Baseline		No Monopoly		No Entry/Exit	
	Effect	P-Value	Effect	P-Value	Effect	P-Value
T+1	-0.102	0.011	-0.093	0.032	-0.059	0.066
T+2	-0.068	0.068	-0.097	0.014	-0.057	0.244
T+3	-0.066	0.076	-0.047	0.228	-0.034	0.408
T+4	-0.077	0.092	-0.031	0.386	-0.047	0.200
T+5	-0.104	0.050	-0.075	0.148	-0.026	0.593
T+6	-0.060	0.267	-0.028	0.666	0.025	0.695
T+7	-0.141	0.086	-0.102	0.257	0.017	0.817
T+8	-0.168	0.019	-0.106	0.169	-0.035	0.655
Treated Routes	10		10		6	

Notes: Data is at the route-quarter level. The dependent variable is the log of the quantity-weighted, route-level average price. The "Baseline" specification uses the control group with no merging firms present to select the five nearest neighbors. The "No Monopoly" specification removes monopoly routes from the control group. The "No Entry/Exit" specification removes treated routes that experienced post-merger entry or exit. "T+X" is the estimate for the Xth quarter after the merger occurred. The effect is the average across all overlap routes. P-values are calculated using the placebo tests detailed in the body of the paper.