The Effects of Jet Fuel Taxes on Aviation Supply Side Outcomes: A Case Study in California

Bachelor Thesis for IBEB
Erasmus University Rotterdam
Erasmus School of Economics

Aniket Godbole
Student Number: 595015

Supervisor: Dr. Dinand Webbink
Second Reviewer: Dr. Josse Delfgaauw

June 2024

Abstract

This paper investigates the effect of aviation taxes on supply side outcomes within the US aviation market. It uses a decrease in the Californian jet-fuel tax from 8.25% to 7.25% in 2011 as a shock to assess how airlines respond regarding average price, passengers transported, number of routes and market share. This fills a gap in existing literature which focuses on climate or employment outcomes by focusing on the microeconomic effects on this salient and large market. We use quarterly airport-level and route-level data from the Bureau of Transportation Statistics to conduct the study. A difference in differences approach is used to assess the impacts of this shock, with several heterogeneity tests being considered. The results suggest that a tax decrease is discharged through the price mechanism, where both the average fare decreases and the passengers served increases. Moreover, market concentration tends to increase. However, in the short-haul market, tax effects are transmitted through the quantity of routes rather than the average fare. These findings may inform policy makers about the understudied responses of the aviation industry to fiscal measures. Given diverse government aims regarding decreasing fares and increasing connectivity, this paper serves an important assessment to one fiscal tool governments may use.

Contents

1 Introduction

Aviation has played an important role in connecting various businesses, people and cultures together, leading to beneficial social and economic outcomes (Cheung et al., 2020). In this vein, its market framework is an important aspect of government policy to facilitate connectivity and low fares that are accessible to a large consumer base. However, with recent developments in environmental studies assessing the impacts of aviation on the environment, such as Ryley et al. (2020), Gossling and Upham (2009), Wuebbles et al. (2007) to name a few, an important balance has emerged for policy outcomes. This balance, with costs to the climate on one side and economic benefits of travel and connectivity on the other side, is regulated through various fiscal and regulatory measures. This paper assesses the economic impacts of one such mechanism, a jet fuel tax, and how it permits the government to exercise control over this economically salient market. A jet fuel tax is a popular way to target airline's most vital commodity, fuel. Unlike usual gasoline, jet fuel has a greater number of additives used to lubricate the engine, and as a result has a very niche usage within commercial flights (MCICO, 2024). This paper assesses a decrease in the jet fuel tax and its effects on airline supply behaviour using a difference in differences approach. Little is known about the ways in which airlines react to changes in taxes, particularly decreases in taxes. This paper fills the gap in existing literature by investigating how airlines respond to the tax by observing effects on the average fare, number of routes, passengers served and market share per route.

The specific tax shock we investigate is a decrease in the jet-fuel tax rate in California. We thus use between-state variation to assess the impact of the intervention on airline supply which adds to existing literature in two crucial ways. Firstly, while many papers consider an increase in aviation taxes, including a jet fuel tax, few analyse effects of the inverse policy change. There is reason to believe the goals of such policies are very different in intended outcomes, and thus may affect supply outcomes differently. A tax decrease might be to encourage cheaper prices and increase connectivity as a primary concern. This paper assesses the viability of this policy instrument. Secondly, prior literature is constrained in it's spread to three key areas. The vast majority of literature focuses on the effects of aviation on environmental outcomes such as carbon emissions, noise pollution or similar climate indicators (Fukui & Miyoshi, 2017; Winchester et al., 2013). A smaller strand of literature studies the effects of tax changes on employment outcomes within airlines such as in Sobieralski and Hubbard (2020). Finally, a significantly smaller strand assesses the effect of the tax on fiscal revenue such as Dama et al. (2023). Therefore, this paper addresses under-analysed outcomes on the supply side of the aviation market. It also considers them within the context of a policy change which lowers the tax, and assesses the intervention's success in four key metrics, elaborated in the Policy Background.

The United States is an ideal setting for such an analysis for three key reasons. Firstly, there is a large availability of data, catalogued by the Bureau of Transportation Statistics (BTS). Secondly, and leading logically to the first, is the fact that the US is a large nation with a high affinity for commercial aviation, with over 650 million passengers according to the BTS. Therefore, this paper's outcome has an important impact vis-a-vis the scale of affected individuals. Thirdly, the US benefits from having a large domestic market where heterogeneity exists within, due to taxation varying at a state level, while also having several airports per state, providing several data points for state-level comparisons.

The paper broadly finds that airlines primarily internalise the tax change through the average fare charged and the number of passengers transported. However, when considering short haul routes, airlines tend to prioritise opening up new routes rather than reducing the fares. These findings pose relevant policy implications for the application of jet fuel taxes in various market settings.

This paper begins with a description of the policy and its normative aims to establish four key variables of interest, and following an assessment of prior literature, formulates hypotheses. The essay then discusses the difference in differences specification used to investigate the causal effect and the relevant data. It concludes with an interpretation and discussion of the results, limitations of this study, and opportunities for future works.

2 Policy Background

California is the largest state by population in the USA and the third largest by area. It has three large hub airports Los Angeles (LAX), San Francisco (SFO), and San Diego (SAN) (F. A. Administration, 2024), the prior two being within the top 20 busiest airports in the US according to the BTS. Furthermore, the state is also geographically separated from the major population centers of the US, thus their residents are frequent users of aviation (See Appendix A). California leads the US aviation market in terms of employment and contribution to the US GDP despite not having the most amount of airports, a statistic outpaced by Texas (Division, 2019).

The specific policy this paper addresses is a decrease in the tax on jet fuel from 8.25% to 7.25% at the end of the second quarter in 2011 in the state of California. This is a state-wide isolated intervention which implies that it doesn't affect states other than California. This tax is applied on fuel on-loaded to the plane that does not originate from the state of California. For instance if there is a flight from New York to California, that lands in California with some fuel remaining; only the additional added fuel for the next flight is taxed under the new regime. This poses two interesting issues. The first overarching note is that this tax regime, protected under Article 24 of the ICAO Chicago Convention in 1944, represents an effective import tax waiver for international flights. This could encourage problems of "tankering", where airlines carry surplus fuel from low tax jurisdictions (Hubert et al., 2015). This however is unlikely to cause large problems in the short run for two reasons. Firstly, supply chain reorganization takes a significant amount of time, particularly for large networks such as airlines. Secondly, the tax differences between California and other states are unlikely to be large enough, owing to similar regulatory bodies, purchasing power and federal authorities; to cause asymmetric tankering from both outside the nation and within the nation. The second overarching issue is that, conversely; the tax might be applied on some fuel not used on flights originating from Californian airports. To resolve the latter issue; the tax authorities state this with regards to who owes jet fuel taxes, *"The aircraft jet fuel tax is owed by the person who owns the aircraft jet fuel when a taxable event occurs (the tax is assessed). Aircraft jet fuel dealers collect the aircraft jet fuel tax for each gallon of aircraft jet fuel sold to an aircraft jet fuel user, delivered into the fuel tank of an aircraft, or into a storage facility from which the fuel is withdrawn for use in an aircraft."* Broadly, the "leftover" fuel is insignificant enough due to modern flight management systems based on various variables such as weather, flight plan and aircraft model/age are able to calculate the necessary fuel required. Therefore, it is assumed not to pose an issue in causal assessment.

The tax decrease is quite a significant 12.12% relative change in taxes and thus is likely to have important outcomes on the airline industry. Jet fuel comprises an important part of the operating expenses of a particular airline, approximately 20.8 percent on average within the US in 2023 (Airlines for America, 2024). Considering that the majority of the aircraft that are commercially profitable are jet fuelled, we may consider jet fuel as synonymous with all airlines fuel costs. Given the high price of jet fuel and it's commonality in commercial aviation, a percentage point decrease is a considerable policy change that should elicit change in consumer outcomes. A crucial aspect is the in-elasticity of airline demand for the commodity. Since effectively all aircraft large enough to be commercially viable are powered by jet engines, it is an irreplaceable component of the industry across the market.

While there is no explicit reasoning provided as to the goals of this policy, we may infer three key goals. The first overarching reason may be the restoration of air demand following the 2008 financial crisis, which put several airlines in problematic economic situations (Scowell, 2012). A direct drop in consumer demand due to loss in finances and often employment, negatively affected both leisure and business airline travel. Furthermore, lack of investment confidence affected equity inflow with stock prices plummeting. This effect raised risk premium on insurance companies which drove up costs for airlines, increasing rates of bankruptcy. Therefore, an assessment on the market share of airlines to see if concentration is dissipating, or incumbents maintain their stronghold over the industry is informative. Similarly, the number of passengers may be an important outcome to consider when analysing a recovering industry to see if demand is stimulated. The second overarching reason for a decrease in taxes is an increasing price trend of jet fuel throughout the time the policy was implemented (U. E. I. Administration, 2024). This increase may have been sufficient for the Californian government, which geographically relies significantly on air travel, to deem it as harming connectivity outcomes. The final reason, may simply be to encourage more tourists to visit California by lowering the cost of travel . This would warrant a study into connectivity improvements, to see if airlines opt into opening up new routes, perhaps in abstract, or at the zero-sum cost of routes to other states. Furthermore, a robust assessment of the average fare of flights could provide insight to the meaningfulness of the tax to affect the most visible consumer outcome, and the primary determinant of aviation demand (Zachariah et al., 2023).

Therefore, the research question then being analysed is as follows;

What are the effects of a decrease in the jet fuel tax rate in California on commercial aviation supply side outcomes?

3 Literature Review

The literature review is divided into four key sections, each discussing previous literature on the relevant outcome variables. As a prelude, there are several studies assessing periphery impacts such as carbon emissions and pollution measures. This is sensible given the primary concern in environmental economics is policy interventions affecting climate, however often papers such as Fukui and Miyoshi (2017) or Winchester et al. (2013) ignore the economic effects of tax measures. A second important comment on prior literature is that while various point of entry along the supply chain for policy intervention exist, the majority of past works consider aviation taxes in a broad sense. They do not distinguish heavily between taxes on fuel or tickets or aircraft and create comparisons. Rather, most literature assumes that taxes in different areas of the supply chain affects the relevant outcome variables similarly. This essay follows a similar approach, and while justifications for using a jet fuel tax to assess supply side outcomes has been provided above, there is no comparison with other mechanisms of government intervention. Rather this essay focuses on assessing if this particular mechanism is valuable to achieve the desired policy outcomes.

The most intuitive outcome variables to discuss when considering tax changes is price and quantity, in this case under the nomenclature of airfares and passengers flown. Standard economic theory in competitive markets would suggest that a decrease in the tax would have simultaneous positive effects on number of passengers and a decrease in the price, since it represents a shift of the supply curve to the right. There is little literature confirming this relationship, giving this paper a unique aspect of scientific relevance. Furthermore, the studies which assess similar relationships are all panel data comparisons, while our study assesses the impact of a tax shock. The value is unlikely to stem from the direction of the relationship itself but rather the magnitude and significance of the relationship, since it can help understand elasticity of the relevant market. The most comprehensive prior assessment is in White et al.

(2019), where they assess all points of government intervention affecting the fare of the ticket, and discover through a panel data approach that taxes are over-shifted to consumers. Trivially, fares increase by more than one dollar for every dollar of tax imposed. However they are unable to suggest concrete justifications for this finding due to network complexity. This however suggest a strong positive relationship between tax increase interventions and the average fare. These results are concurred by Chuang (2020) who estimates the over shifting between 30-70 cents for every dollar tax imposed, suggesting that tax salience might be the cause. An over-shifting implies that for a single dollar increase in taxes, the prices increase by more than one dollar. Karlsson et al. (2004) provide additional insight with regard to taxes on infrastructure. They notice that the most intense effect is felt by the cheapest tickets due to the high relative proportion of tax-based costs and thin margins. Therefore the first hypothesis of this paper is;

Hypothesis 1: The intervention decreasing the tax rate on jet fuel will cause a decrease in the average fare in California relative to the rest of the nation

The second hypothesis considers the counterpart of price, quantity. A decrease in the tax implies an increase in the number of passengers as a mechanical result of the decrease in the price. However, given the extremely strong effect of jet fuel taxes on price discussed in previous literature, we may infer the total demand curve is relatively price inelastic compared to supply. This would imply a comparatively smaller effect on the number of passengers. Therefore, while there should be a positive effect on the number of passengers, it may be smaller and thus statistically insignificant. This leads to our second hypothesis:

Hypothesis 2: The intervention decreasing the tax rate on jet fuel will cause an increase in the number of passenger in California relative to the rest of the nation, however it is likely small in magnitude

The third hypothesis focuses on the number of routes served by an airport, also referred to as connectivity. Sobieralski and Hubbard (2020) discusses issue regarding employment, traffic and emissions using a difference in differences (DiD) analysis within the US. Relevant to this particular study, they find a positive short run effect on air traffic and therefore connectivity, however the effect fades within the year as airlines adjust better to the market conditions. This study is worth noting since it is the only study that considers a jet fuel tax decrease in comparison to other papers which focus on a tax increase. Similarly, though within a different geographical setting, Bernardo et al. (2024) considers the effect of an increase in a "flight" tax in Europe. It has a particular focus on low-cost short-haul flights and determines through a staggered DiD analysis that the taxes decrease the average flights per route by 12%. The paper further notices that, given the demand for low cost carriers (LCCs) is relatively elastic, passengers may incur a greater proportion of the tax burden, therefore flights with lower "added value" such as leisure and holiday trips are less likely to persist. This may be due to a lower willingness to pay, but also because short haul flights have reasonable inter-modal alternatives such as buses and trains which make them more price sensitive Commission (2021). While not explicitly considering connectivity, but rather frequency of already existing connections, both these studies indicate an inverse association between the tax and the connectivity of an airline's route structure. Thus, the third hypothesis of this paper is:

Hypothesis 3: The intervention decreasing the tax rate on jet fuel will cause an increase in the number of routes served by airports in California relative to the rest of the nation, however it is likely small in magnitude

The final variable addressed by this paper is the impact of taxes on the market share per route of various airlines. In particular, we assess the effect on the largest airline on a particular city-pair route and the cheapest airline on the same route, the details of which are better addressed in Data. This relationship has a surprisingly small amount of academic literature hitherto, thus we must refer to economic intuition to formulate a hypothesis. Three premises are clear with regards to route-market structure. The first premise is that there is intense monopolistic competition on particular routes. This historically has been driven by several airline mergers (Gil & Kim, 2021) and high barriers of entry to the market itself with regards to start-up costs and spatial first mover advantages such as "hub airports"(airports from which the vast majority of their flights originate) (Lijesen et al., 2002). The second premise is that larger airlines tend to have more flights and thus use more fuel. The third premise is that, on a route level, airlines operate closer to a Cournout oligopoly model as opposed to any other form (Koopmans & Lieshout, 2016), (Lijesen et al., 2002). This model suggests that, on a particular route, airlines compete on number of flights and timings of flights more than price, due to some aspects of product differentiation between airlines. These facts would suggest that airlines who are larger consumers of jet fuel would be more affected by the tax decrease and thus be able to free up capital to operate greater frequencies on a route. Concurrently, airlines with smaller market shares on the route are likely to have opposite reaction and attempt to escape competition on one route, to consolidate market power on other routes. This follows closely in line with standard Cournout oligopoly theory.

A similar intuition applies for the airlines with the cheapest fare. Given the relative price sensitivity of passengers, particularly on low cost carriers, they are likely to consolidate market share on routes by competing on price segment of the market. This would also result in an increase in their market share due to their ability out price-out other competitors by lowering fares further. This may also unlock pre-existing latent demand, who found tickets too expensive *ex-ante*. Furthermore, low cost carriers are particularly known to hedge their fuel costs by purchasing several gallons in advance (Smyth & Pearce, 2006). Therefore, given a tax cut, LCCs are likely to capitalise by purchasing large quantities of fuel and storing it. This allows them greater flexibility in the short run and permits them to further decrease price therefore capturing a greater proportion of the total demand. Furthermore, the set of airlines which are the largest carrier on a particular route and the cheapest carrier might be equivalent due to economies of scale advantages on a route (Johnston & Ozment, 2013).

While there is an argument that market share per route of airlines should decrease due to lower costs for smaller airlines to be able to expand, this is likely not a nuanced perspective. While the costs are lower, they asymmetrically benefit airlines with large pre-existing infrastructure in place such as catering, baggage crew and landing slots at airports. Strategically, it makes little sense for smaller airlines to compete on a route where legacy carriers have established brands. In totality, this information leads to our fourth and final hypothesis:

Hypothesis 4: The intervention decreasing the tax rate on jet fuel will cause an increase in the market share of the largest as well as the cheapest airline on a particular route in California relative to the rest of the nation

4 Methodology

To test this research question we run a series of DiD regressions on separate outcome variables analysing the effect of three separate policies. The most rudimentary DiD regression equation is specified below;

$$Y_{it} = \alpha_{it} + \gamma post_t + \rho CA_i + \beta DID_{it} + \epsilon_{it} \tag{1}$$

where:

$$DID_{it} = post_t * CA_i \tag{2}$$

We use a standard Ordinary Least Squares linear regression to estimate this causal effect. The outcome variables are the four variables of interest specified within each hypothesis; average fare for an airport, number of routes served by an airport, number of passengers served by an airport and market share of the largest/cheapest airline on a particular route. The first three outcome variables can be seen in the airport-level data, and the latter outcome variable may be seen in the route-level data. The variable **post** is a dummy variable indicating if a particular time period is before the intervention (post = 0), or after the intervention (post=1). The variable **CA** is a dummy variable indicating if a particular observation is in the control group (CA=0) or in the treatment group California (CA=1). The variable **DID** represents the treatment variable and is a binary variable. It takes the value of 1 when an observation was treated in a particular time period and a value of 0 when the observation was not treated in a particular time period. The treatment dummy can also be seen as an interaction terms between a treat group dummy and a treatment period dummy, seen in equation 2. The β is our Difference in Difference estimator, henceforth called the Average Treatment Effect on the Treated (ATET). The parameter α is the value of the outcome variable for the control group prior to the intervention, and acts as a constant within the linear regression and is a unit fixed effect within the context of the DiD regression. The parameter ϵ is the error term and considers all variation not explained by the model.

This paper, in order to control for state and temporal heterogeneity, will include time and group fixed effects as follows:

$$Y_{it} = \alpha_{it} + \gamma post_t + \rho CA_i + \beta DID_{it} + \sum_{j=1}^{46} \theta_j \text{state}_j + \sum_{k=1}^{112} \lambda_k \text{time}_k + \epsilon_{it} \tag{3}$$

where θ considers the group-fixed effects by being multiplied by a state binary for all 46 states within the dataset. Similarly, λ considers the time-fixed effects by being multiplied by each of the 112 quarters between January 1996 to December 2023. The variable for California will be dropped from the state-fixed effects due to issues of perfect multi-colinearity with the treatment group variable.

4.1 Standard Errors

For all regressions henceforth, state cluster standard errors are used, since treatment is applied at a state level, i.e. for all airports within California. Observations may be subdivided into smaller clusters, where treatment assignment is correlated within each cluster. Therefore, standard heteroskedasticity robust standard errors would be inappropriate since airport outcomes between each state are not independently distributed. By clustering standard errors at the state level, we remove this source of heterogeneity.

4.2 Verifying assumptions

There are two key assumptions that must hold for a valid causal estimate within the DiD framework. The first is an absence of confounding events during the intervention taking place place. If another major shock occurs, it is unclear if changes in averages of the outcome variable are caused by the intervention studied or exogenous factors. Fortunately for this paper, we have quarterly data, which means the time frame is more precise and the probability of an intervention/shock asymmetrically affecting the treatment and control group is small. Furthermore, according to CDTFA, 2024, there has been no change of the Sales Tax rate since 2010, which is the extent of data publicly available. The only small discrepancy is a 20 cent/gallon increase in the ”Prepayment of Sales Tax Rate”. This theoretically affects the airlines since they receive less money back in tax overpay reimbursement. However, since this is a yearly tax, and is contingent on already purchased jet fuel, it is unlikely to pose a large issue. Furthermore, the tax increase is quite small and unit based rather than percentage-based, and thus not likely to influence airline decision making as much. Finally, the increase in tax is in the opposite direction to the tax being assessed in this paper. In other words it is an increase in the tax; which underestimates the causal effect or may indicate that there is no effect when there actually is one. Thus we assess the lower bound of the effect.

The second key assumption is the parallel trends assumption. This assumption states that if the treatment was not applied, the control and treatment group would have equivalent trends. We first look at trend plots of all the variables and see if they are visually parallel. These graphs may be seen in Appendix B. For a statistical assessment, we use a test of introducing lead variables into the DiD regression. A "lead" is a variable that artificially moves the treatment one period ex-ante of the actual treatment. Therefore, when inputting leads into the DiD regressions, we create a placebo test to see if there is a confounding shock occurring prior to the real intervention that masks the causal effect. The leads chosen are the first lead, which indicates the quarter prior to the policy implementation, and the fourth lead, which indicates the analogous quarter in the year prior to the intervention. In this case the two chosen leads would represent 2011Q1 (first lead) and 2010Q2 (fourth lead). There are very severe limitations of the leads method of testing parallel trends. The most important one is that it is extremely sensitive to the model specification. Several different choices of leads, combinations of controls and methodology of standard error calculation have both artificially confirmed and rejected the assumption. Therefore, we prioritise the visual representation of the parallel trends than the highly volatile lead method.

This test becomes slightly more complicated when considering route-level data. The formatting of this dataset is such that each observation is a set of directionless city pairs. This means that when considering a route, let us assume for example, from Los Angeles (city1) to New York JFK (city2), the data considers all the bidirectional flights between the cities. Therefore, when considering group fixed effects, it becomes difficult to isolate the departing airport and the destination airport. We then run a following DiD regression:

$$Y_{it} = \alpha_{it} + \gamma post_t + \rho CA_i + \sum_{m=0}^{2} \beta_m \mathrm{DiD}_{i,t+m} + \sum_{j=1}^{46} \theta_j \mathrm{state}_j + \sum_{k=1}^{112} \lambda_k \mathrm{time}_k + \epsilon_{it} \tag{4}$$

where the variable DiD is the variable indicating if the treatment took place in particular period. We include two lags to test if parallel trends hold for the variables. However, since the group variable for state is not isolated but rather split through directionless route pairs, we must test the above regression considering both group variables for state. Therefore, we run two regressions, each one considering a different group (city1 or city2) for group fixed effects. Only if both regressions yield insignificant leads may we assume parallel trends.

5 Data

The data used is derived from the US Department of Transportation, specifically the Bureau of Transportation Statistics (Keizer, 2024a)(Keizer, 2024b). The Domestic Consumer Airfare report is a vast panel dataset using quarterly data on a several consumer outcomes during the time period 1996 - 2023. Two particular datasets are of interest, containing route-level data and airport-level data. The route-level data contains the top 1000 US largest city market pairs in the contiguous United States (excluding Alaska and Hawaii) (Keizer, 2024a). Every observation records one route in a particular quarter and has information on the average market fare, the market share of the largest and lowest fare airline. The dataset contains other variables however they are not relevant for this paper to consider. The summary statistics for route-level data can be seen in Table 1. It is important to consider that the number of observations for the the market share per route of the largest airline is lower than for the cheapest airline. These observations are all missing for the cheapest carrier in 1996 and 1997 for routes from Aspen, Colorado to both Chicago and New York. These three observations seem to be randomly missing, but are unlikely to considerably affect the internal validity the estimates given that Aspen itself has 61 other data points both prior and after the intervention.

Table 1: Summary Statistics of Route Level Data

Variable	**Control** (1)	**Treated** (2)
Market Share (Largest Airline)	0.56 [0.18]	0.52 [0.20]
N	105,462	6,572
Market Share (Cheapest Airline)	0.32 [0.24]	0.32 [0.24]
N	105,459	6,572

Notes: The descriptive statistics for the market share from route-level data.. Column 1 indicates means for the control group and Column 2 indicates means for the treatment group (California). The means consider the entire period for which data is available; 1996Q1 - 2023Q4. Standard deviations are shown in square brackets. The number of observations is denoted by N. Market share is represented as a proportion.

Airport-level data records data classified by airport instead of route for all airports with flights that carry more than 20 passengers per day (Keizer, 2024b). It records the number of routes served by an airport, passengers passing through the airport, the average fare and the premium. The summary statistics can be seen in Table 2. There are again discrepancies in the number of observations between the total, short haul and long haul specifications. These are due to the lack of either data availability, or the lack of flights that fit within each specification. For instance, "Martha's Vineyard" is an airport in Massachusetts which only serves short shuttle flights from the island of Martha's Vineyard to airports nearby. Therefore, it has no data for long haul routes. A similar logic is applied for airports far away from their closest destinations, and thus having only long haul flights. This may change as airports gain new routes. Another reason for missing observations might be that certain smaller airports do not report their data in some years. This, while not explicitly seen in the number of observations, creates some inconsistencies. However, as with route-level data, these omissions are unlikely to cause significant problems for our estimates due to few missing observations relative to the total.

We operationalise four key outcome variables. The first variable is the average fare. This proxies how sensitive airlines are to changes in tax policy through the price mechanism. The rationale for this is two fold. Firstly, there are more airports considered within the airport-level data, since airports only need a flight carrying 20 passengers. This allows for a larger number of observations and a greater statistical power of our DiD regressions. Secondly, the route-level data considers a set of directionless city market pairs. This is problematic because the incidence of the tax is only in the city where the route originates and thus it is hard to isolate where the tax change is imposed. Therefore, we structurally underestimate our causal estimate and the significance of them. Thus it is important to highlight, whenever route-level data is used, we estimate the lower bound of the causal effect.

The second outcome variable is the number of routes served by an airport. For this we can only use airport-level data since it is the only table that has the data for markets served by an airport. It calculates the number of other airports that the relevant airport serves with flights of over 20 passengers per day for a particular quarter.

The third outcome variable is passengers flown. For the same reasons as the average fare, we use airport-level for this data. This variable records the directionless (departures and arrivals) flow of passengers passing through a particular airport for a particular quarter. The paper will also consider the effect on the natural logarithm of passengers. Seen in Table 2, the minimum number of flights in the control group is 120, while the mean is over 800,000. This implies a large left skewed distribution, a pattern that persists when considering both short and long haul passenger data as well. Therefore, to

Table 2: Summary Statistics of Airport Level Data

Variable	**Control** (1)	**Treated** (2)
Routes (Total)	62.97	104.89
	[47.72]	[65.57]
Passengers (Total)	829,967.60	1,613,463.00
	[1,406,726.00]	[1,950,555.00]
Average Fare (Total)	206.67	196.02
	[50.20]	[43.39]
N	20,292	1,727
Routes (Short Haul)	22.72	17.26
	[19.69]	[8.90]
Passengers (Short Haul)	300,627.50	673,495.10
	[504,296.20]	[638,263.90]
Average Fare (Short Haul)	188.59	136.07
	[52.69]	[36.33]
N	20,000	1,721
Routes (Long Haul)	41.10	89.51
	[32.45]	[56.18]
Passengers (Long Haul)	540,592.70	961,800.20
	[995,542.50]	[1,427,809.00]
Average Fare (Long Haul)	227.45	253.17
	[60.66]	[50.60]
N	20,032	1,692
Airports	190	16

Notes: The means for important outcome variables seen in airport-level data are seen in Column 1 (for control group) and Column 2 (for California). Means are calculated across all 112 quarters considered in the dataset; between 1996Q1 and 2023Q4. The standard deviations are in square brackets. The number of observations are stated for different levels of heterogeneity. The fare variables are all measured in US Dollars at the time of recording. The number of observations is denoted by N. Passengers and Routes are both count variables. Short haul only and long haul only flights are also considered.

preserve ordinality but mitigate the effects of outlier, we also consider the natural logarithm as a form of sensitivity analysis.

The fourth and final outcome variable is the market share per route, which may be found in route-level data. The relevant heterogeneity variables for largest airlines per route and cheapest airlines per route is also present in the data. It records the market share of the largest airline and the airline with the lowest fare on a particular route, in percentages, for a particular quarter. This estimates the lower bound of the causal effect, since it sometimes assumes that the tax intervention was applied where it is not. While the magnitude of the causal estimate is likely to be underestimated, we may still be able to meaningfully infer causality in terms of the direction of the effect.

6 Results

The results section starts with a preliminary assessment of the common trends assumption, a prerequisite to the internal validity of the DiD method. It then discusses both airport-level data and route-level data separately. Finally, it concludes with a sensitivity analysis to validate the robustness of the primary results to subsets of the data.

6.1 Parallel Trends

In Table 3, we see a testing of the parallel trends assumption for the relevant outcome variables in the airport-level data. By inputting placebo treatments *ex-ante* the actual treatment, we test if artificially inserting a treatment shows significant effects. This would indicate that the trends are not parallel and the DiD estimate is invalid. We see however, that there are certain leads that are significant at a 1% level. However, as discussed above in Methodology, the leads method is very prone to specification manipulation. To make sure that parallel trends do not hold in the instances where the leads method fails, we seek visual confirmation. These graphs may be seen in Appendix B: Parallel Trends, where on the left plots we see the observed means of the variable of interest, and on the right plots we see a linear trends model. The linear trends model attempts to fit the treatment group means to the control group means considering the trends and standard errors of the treatment group across time. A more in depth explanation of how this is done can be found in Stata's manual of the didregress command. The figures 4 to 11 display the means in each time period for both the treatment and control group. We can see that the majority of the trends are indeed parallel and thus we may trust the DiD estimator. The only possible issue arises when considering the outcome variable of the long haul log(passengers), which seems to have diverging means earlier than the intervention is enacted. While we still estimate the regression, the coefficient will not be interpreted due to the failure of the parallel trends assumption.

In Table 4 we are testing the parallel trends assumption for the market share using route-level data. The aforementioned problems with regards to group fixed effects leads us to run two regressions for each outcome variable. Each choice of group also indicates the group used to cluster the standard errors. In both regressions considering the lowest fare airline on a particular route (column 3 and 4), we see both the first and second leads are insignificant. Therefore, at a 10% level, we may proceed with our DiD estimations for the market share per route of the cheapest airline. However, the market share of the largest airline on a particular route poses issues with common trends. The variable "lead1" is significant in one of the regressions. This poses some concerns with regards to the internal validity of the DiD estimate. We will still cautiously interpret the ATET for an associative approximation, however stress that further and more robust research should be conducted, perhaps in different settings to test this effect more accurately.

Table 3: Common Trends Test Using the Leads Method for Airport Level Data

Panel A: Total				
Variable	Fare (1)	Routes (2)	Passengers (3)	Log(Passengers) (4)
DiD	-6.850**	-3.618**	24397.343	-0.031
	(2.704)	(1.395)	(26184.483)	(0.056)
t	76.025***	11.322***	606095.940***	0.0.898***
	(7.254)	(2.415)	(73855.874)	(0.114)
CA	-46.729***	60.327***	1.381e+06***	1.654***
	(1.483)	(0.547)	(10215.261)	(0.022)
lead1	0.133	-3.301***	-17602.293	-0.095
	(1.465)	(1.040)	(16914.113)	(0.051)
lead4	0.740	7.639***	68646.267***	0.090
	(2.581)	(1.205)	(21582.251)	(0.054)
_cons	215.230***	36.800***	-1.138e+05***	11.125
	(4.286)	(1.068)	(38173.453)	(0.065)
N	21819	21819	21819	21819
r^2	0.425	0.286	0.179	0.192

Panel B: Short Haul				
Variable	Fare (1)	Routes (2)	Passengers (3)	Log(Passengers) (4)
DiD	0.377	1.401***	22334.437**	0.066
	(2.726)	(0.479)	(8668.529)	(0.062)
t	70.675***	0.346	144384.452***	0.654***
	(7.917)	(0.969)	(25668.257)	(0.126)
CA	-94.528***	-3.661***	552394.207***	1.377***
	(1.932)	(0.267)	(4684.268)	(0.029)
lead1	-0.893	-1.209***	-34947.551***	-0.146**
	(1.694)	(0.395)	(6081.392)	(0.063)
lead4	4.087	1.699***	13548.454	0.161**
	(3.849)	(0.456)	(9165.772)	(0.068)
_cons	195.956	19.640***	34806.917*	10.688***
	(4.512)	(0.539)	(17491.996)	(0.067)
N	21522	21522	21522	21522
r^2	0.517	0.371	0.241	0.222

Panel C: Long Haul				
Variable	Fare (1)	Routes (2)	Passengers (3)	Log(Passengers) (4)
DiD	-0.809	-3.911***	15098.594	-0.173***
	(3.039)	(0.964)	(19235.017)	(0.059)
t	76.068***	9.402***	450485.895***	1.011***
	(7.436)	(1.747)	(63683.413)	(0.123)
CA	-16.487***	66.193***	855410.148***	2.013***
	(1.482)	(0.335)	(8962.578)	(0.023)
lead1	-3.196*	-1.823**	19750.068	-0.038
	(1.768)	(0.766)	(12162.029)	(0.057)
lead4	-3.435	3.788***	29580.753*	-0.060
	(2.281)	(0.829)	(15725.479)	(0.060)
_cons	248.177***	17.959***	-1.465e+05***	10.141***
	(4.266)	(0.866)	(29515.934)	(0.071)
N	21528	21528	21528	21528
r^2	0.439	0.357	0.166	0.211

Notes: Common trends tests using the leads method. Time and Group fixed effects are included in each regression, however, the relevant dummies are omitted for presentation purposes. The DiD variable is the treatment indicator and the prefix "lead" indicates the relevant placebo treatments prior to the real treatment. The sample size, denoted by N, and the r-squared are displayed at the bottom of each panel. Stars indicate significance where * p≤0.1, ** p≤0.05, *** p≤0.01.

Table 4: Common Trends Test Using the Leads Method for Route Level Data

Variable	Largest Airline 1 (1)	Largest Airline 2 (2)	Cheapest Airline 1 (3)	Cheapest Airline 2 (4)
DID	0.022***	0.024***	0.033***	0.035***
	(0.008)	(0.009)	(0.013)	(0.013)
CA	-0.100***	-0.090**	-0.083***	-0.049
	(0.024)	(0.043)	(0.017)	(0.035)
t	-0.014	-0.012	0.024	0.030*
	(0.011)	(0.017)	(0.016)	(0.016)
lead1	-0.006*	-0.005	0.008	0.009
	(0.004)	(0.004)	(0.009)	(0.006)
lead4	0.009	0.007	-0.013	-0.013
	(0.007)	(0.012)	(0.014)	(0.010)
cons	0.610***	0.929***	0.469***	0.690***
	(0.011)	(0.006)	(0.012)	(0.010)
N	112034	112034	112031	112031
r^2	0.147	0.090	0.088	0.061

Notes: Common trends tests using the leads method. Time and Group fixed effects are included in each regression, however, the relevant dummies are omitted for presentation purposes. The DiD variable is the treatment indicator and the prefix "lead" indicates the relevant placebo treatments prior to the real treatment. The sample size is denoted by N, and the r-squared is denoted by r^2. Stars indicate significance where * $p \leq 0.1$, ** $p \leq 0.05$, *** $p \leq 0.01$.

6.2 Primary Airport Level Results

Panel A of Table 5 shows us the DiD estimates for the effect of the policy intervention when considering the market as a whole. The average treatment effect of the intervention on the average fare (seen in Column 1) is negative and significant at a 10% level. It indicates that a one percentage point decrease in the tax reduces fares by just over six dollars. This represents almost a 3% relative drop in the average fare, which is a very economically significant finding. The direction of the effect is in accordance with a standard competitive market where a tax decrease causes a decrease in the price consumers pay. However the magnitude of the coefficient indicates that this is a rather large effect, concurring with White et al. (2019) and Chuang (2020). This indicates that the over shifting of the tax works inversely as well, as prior research has only considered tax increases. This has important policy implications considering the governments ability to affect the fares consumers have to pay. It indicates that a jet fuel tax is a powerful policy tool that governments can use, where even small changes have significant effects on airline responses.

The average treatment effect of the intervention on the number of routes can be seen in Panel A Column 2. The coefficient is positive however insignificant at a 10% level, thus we cannot reject the null hypothesis that the coefficient is equal to zero. This implies that the number of routes is not statistically affected by a one percentage point decrease in the taxes. This might be explained by viewing Panel B and Panel C for the number of routes briefly. We see in Panel B that there is a statistically significant increase in the number of routes served and in Panel C that there is a statistically significant decrease in the number of routes. The reasons for this will be discussed later in the essay. On net however, a tax decrease on jet fuel does not seem to be an effective policy tool to raise connectivity.

Finally we discuss both passenger outcome variables, seen in column 3 and 4 of Panel A in Table 5. While there is no statistically significant effect on the natural logarithm of passengers, we see a positive coefficient in the absolute number of passengers served which is statistically significant at a 1% level. This further reinforces the findings we derive with regards to the fare and is consistent with a competitive market framework. A decrease in the tax, in a perfectly competitive supply-demand framework, shifts

Table 5: Difference-in-Differences Regression Results for Airport Level Data

Panel A: Total				
Variable	Fare (1)	Routes (2)	Passengers (3)	ln(Passengers) (4)
ATET	-6.026*	0.283	71311.158***	-0.040
	(3.270)	(1.159)	(22580.325)	(0.049)
cons	176.351***	59.077***	582424.205***	11.854***
	(4.204)	(1.042)	(37859.182)	(0.064)
Mean (Control)	206.669	62.974	829967.600	13.629
N	21819	21819	21819	21819
Panel B: Short Haul				
Variable	Fare (1)	Routes (2)	Passengers (3)	ln(Passengers) (4)
ATET	3.321	1.801***	641.149	0.073
	(4.122)	(0.580)	(10161.124)	(0.063)
cons	151.687***	21.816***	245143.718***	11.050***
	(4.427)	(0.528)	(17434.005)	(0.066)
Mean (Control)	188.594	22.724	300627.500	12.614
N	21522	21522	21522	21522
Panel C: Long Haul				
Variable	Fare (1)	Routes (2)	Passengers (3)	ln(Passengers) (4)
ATET	-7.159**	-2.165***	62144.443***	-0.266***
	(3.335)	(0.708)	(20101.694)	(0.048)
cons	205.325***	38.892***	350790.524***	11.191***
	(4.212)	(0.863)	(29082.579)	(0.071)
Mean (Control)	227.448	41.104	540592.700	13.200
N	21528	21528	21528	21528

Notes: Difference in difference regression estimation with the respective outcome variables. Time and Group fixed effects and interaction terms are included in each regression, however, the respective dummies are omitted for presentation purposes. The ATET is the coefficient of the treatment variable dummy which indicates if an observation was treated in a particular time period. The means of the control group are also provided for a better interpretation of the magnitude of the ATET by considering relative effects. Standard errors were clustered at the state-level. The sample size is denoted by N. Stars indicate significance where * $p \leq 0.1$, ** $p \leq 0.05$, *** $p \leq 0.01$.

the supply curve to the right, thus lowering price and raising quantity. In this case, price is analogous to the average fare that consumers pay and the quantity is the number of passengers served as the product airlines sell is a seat on an airplane. The increase is economically quite large, representing an 8% relative increase over the control group, indicating large policy effectiveness.

On net, tax effects seem to be large and transmitted via the fare of the tickets rather than changes in connectivity. It concurs with both prior literature and standard economic theory about a perfectly competitive market. Therefore, our first and second hypothesis likely hold.However, there seems to be no statistically significant effect on the number of routes served, in contradiction with the third hypothesis. The magnitude of the results seem to indicate that tax changes and decisions regarding jet fuel pose and important economic effect for the broad aviation market.

6.3 Heterogeneity Analysis

6.3.1 Short Haul Flights

In Panel B of Table 5 we see the effect of the policy intervention on the relevant outcome variables when considering only short haul flights. The average treatment effect in columns 1, 3 and 4 are all insignificant at a 10% level. The contrasts from the results when considering all flights quite considerably. Most notably, the transmission of the tax does not seem to be via the standard mechanism of price decreases.

Instead, seen in Column 2 is the average treatment effect on the number of routes operating from airports. The coefficient is positive and significant at a 1% level. It indicates that, on average, a one percentage point decrease in the tax increases the number of routes by 1.8. This is an 8% relative increase, which presents an economically large and significant finding. It strongly suggests that when segmenting the market, the tax effects are transmitted through an increase in connectivity with airlines using the opportunity to open new routes, instead of reducing prices on already existing routes.

The most compelling reason for a short-haul increase in connectivity is the demand elasticity of the market. Passengers may have increased their demand due to a mechanical decrease in the price of the ticket, owing to the tax decrease. The latent demands for routes that were prior cost inefficient, now became financially viable under the new tax regime. Bernardo et al. (2024) discusses in depth that short haul and particularly low cost carriers serve markets with a greater price elasticity due to a larger proportion of "non-essential travel" such as leisure trips and holidays. Furthermore, as mentioned by Commission (2021), short haul flights experience inter-modal competition such as trains, buses and cars, which causes its consumer base to be more price-sensitive. This may explain the relatively large change in the connectivity. Furthermore, short-haul routes may be more appealing to open, particularly in contrast to long-haul routes given many short-haul flights are domestic flights. This may pose asymmetric benefits regarding regulatory costs, since international flights are more procedural intensive. Comparatively, long-haul international flights are planned well in advance and may suffer from lack of flexibility.

This effect on short haul routes is important as it indicates that a jet-fuel tax decrease is an effective mechanism to increase connectivity between short haul destinations. These results provide a more nuanced understanding of the second hypothesis. Since opening up new routes is a large initial cost, there is only an effect of the tax when it becomes financially viable, which is more likely in the short haul market segment due to high demand elasticity. This may have relevant policy implications in terms of domestic routes to smaller cities that were prior deemed economically irrelevant for airlines. Further, it may have important inequality decreasing outcomes, where lower-income areas are more likely to be served with greater connectivity. It also implicitly suggests that airlines prioritise starting new routes over only increasing frequency on already existing routes. However, in terms of improving access in already present routes, a jet fuel tax seems to be ineffective within the short haul market.

6.3.2 Long Haul Flights

In Panel C of Table of Table 5 we see the effect of the intervention on relevant outcome variables when considering only long haul flights. In column 1 we see the average treatment effect of the intervention on the fare of the flights. The coefficient is negative and significant at a 5% significance level. It suggests that, on average, a one percentage point decrease in the tax on jet fuel translates into a seven dollar decrease in the average fare. This coefficient is notably larger than the coefficient when considering total flights, likely due to the fact that more fuel is used on long-haul flights, thus the tax effects are stronger. Furthermore, the effect is in line with both standard competitive market theory and prior literature. The seven dollar increase represents a similar 3% increase in the fare, which is quite economically large

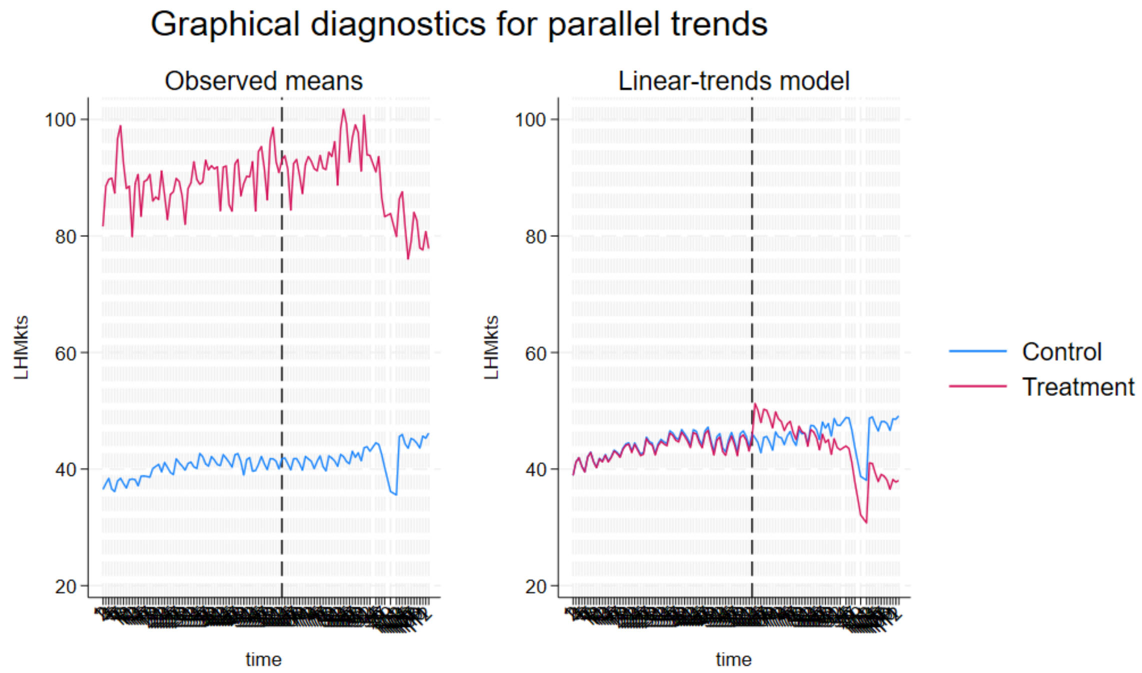

Figure 1: Trends for the Long Haul Number of Routes

and significant. It indicates that a jet fuel tax decrease is a strong policy mechanism to affect prices on long haul routes in particular.

In column 2 we see the average treatment effect on the number of long haul routes. The coefficient is negative and statistically significant at a 1% level. It suggests that, on average, the intervention to decrease the jet fuel tax by one percent results in approximately two routes fewer than the control group. This is a rather counter-intuitive finding, which is more accurately understood by viewing the trend graphically (seen in Figure 1).

Figure 1 shows on the left plot, the observed means comparison between the control and treatment group. The right plot, displays a linear trend model of the treatment group over the control group, making the two groups easier to compare. It indicates that while the average treatment effect may be negative, in the short run, we see the number of routes increase. However this effect quickly decays in the long-run. This nuance explains the significant negative coefficient, which may not be the real effect of the tax decrease.

In column 3 and 4. the ATET on both passenger outcome variables are displayed. As discussed in Parallel Trends, we do not interpret the natural logarithm of passengers due to it's failure to comply with the parallel trends assumption. Instead we interpret the coefficient in column 3 which is positive and significant at a 1% level. It suggests that, on average, the effect decreasing the jet fuel tax rate by 1% increases the number of passengers served by just over sixty thousand people. This represents a nearly 11% relative increase, which is a very economically significant finding. It further concurs with the standard competitive market framework of a tax decrease, which simultaneously decreases prices and increases quantity provided.

In sum, the jet fuel tax decrease seems to have the largest effects on the long haul segment of the market. This is rather intuitive given long-haul flights use a larger amount of fuel than average and thus are more directly affected by changes to tax policy. The results also suggest a competitive market structure where the effects of the tax are discharged via the fare airlines charge, and the passengers they serve. It has ambiguous effects on the number of routes served, and this outcome variable would benefit from further statistical analyses.

6.4 Primary Route Level Results

In Table 6 , we see the results of the DiD regressions for the market shares of the largest airline and cheapest airline on every route, with time and group fixed effects included. In column 1 and 2 we see the regression estimates for the effect of the intervention on the market share of the largest airline. At the lower bound of significance, with a 5% significance level, we may reject the null hypothesis that the intervention has no effect on the market share of the largest airline. The regression would suggest that a percentage point decrease in the tax on aviation jet fuel would on average increase the market share by a range between 0.025% and 0.026%. At the lower bound, this represents approximately a 4% relative increase in the market share, an economically significant effect which may indicate competition concerns of a tax decrease. This concurs with the fourth hypothesis, where larger airlines consolidate power on routes they already have a competitive advantage within.

In column 3 and 4, the regression estimates for the market share of the cheapest airline on a route are shown. With both state fixed effect and cluster specifications, the coefficients are significant at a 10% level at the lower bound and a 1% level at the upper bound. We see a positive coefficient in both regression specifications with the lower bound estimate indicating that a 1 percentage point decrease in the tax, on average, increases the market share by 2.9%. This represents an equally economically significant 5% relative increase in the market share of the low fare airline, raising similar competition concerns. These results also concur with the fourth hypothesis.

By imposing a tax decrease, governments risk increasing market concentration and thus enhance oligopoly benefits. This may not be inherently problematic, since large airlines benefit consumers through more comprehensively integrated route networks, allowing ease of connectivity. However it prevents innovation within the industry by allowing new entrants, and may increase price markups and inefficiencies in the long run when smaller airlines are priced out.

Table 6: Difference in Differences Regression for Route Level Data

Variable	Largest Airline 1 (1)	Largest Airline 2 (2)	Cheapest Airline 1 (3)	Cheapest Airline 2 (4)
DID	0.025***	0.026**	0.029*	0.032***
	(0.007)	(0.010)	(0.015)	(0.011)
CA	-0.100***	-0.089**	-0.084***	-0.050
	(0.024)	(0.043)	(0.017)	(0.035)
t	-0.014	-0.012	0.024	0.030*
	(0.011)	(0.017)	(0.015)	(0.016)
cons	0.610***	0.929***	0.469***	0.690***
	(0.011)	(0.006)	(0.011)	(0.010)
Mean (Control)	0.56	0.56	0.52	0.52
r^2	0.147	0.090	0.088	0.061
N	112034	112034	112031	112031

Notes: Difference in difference regression estimation. Time and Group fixed effects are included in each regression however the respective dummies in interest of condensed presentation are hidden from the table. Cluster standard errors were used corresponding to the state group chosen for the fixed effects estimates. The sample size, denoted by 'N', and the r-squared are displayed at the bottom of the table. Stars indicate significance where * p≤0.1, ** p≤0.05, *** p≤0.01.

7 Sensitivity Tests

In order to test the robustness of the results, we test it's validity under a subset of the full data. To this end, we filter the data to only include the 20 most populous states, excluding California (21 total states including California). This is a reasonable subset to consider for both logical and statistical reasons. California is the largest state in the US and thus might be very different qualitatively than smaller states with lesser demand. Therefore, indicators such as average fare or connectivity are likely to differ significantly. This may bias our estimates in several ways. For example, since the demand is lower in smaller states, they may be more affected by spillover effects of the policy intervention. Therefore, the opposite effect such as a fare increase or connectivity decrease, due to California's increased competitiveness, is likely to affect smaller states more due to their lack of gravity or competitive advantages. Therefore, larger states are more likely to be similar to California with regards to their aviation market and thus justify their place in the subset. Statistically, the top 20 provide the additional advantage of preserving enough data points to make the regression meaningful. A very small set of data points may lead to statistically insignificant results due to low statistical power. Therefore, the 20 control group states considered are: Texas, Florida, New York, Pennsylvania, Illinois, Ohio, Georgia, North Carolina, Michigan, New Jersey, Virginia, Washington, Arizona, Massachusetts, Tennessee, Indiana, Missouri, Maryland, Wisconsin, and Colorado. Therefore, this is an statistically and economically relevant counterfactual subset to study.

Notably, we can only conduct sensitivity tests meaningfully on airport-level data, since the route-level data is organised such that there are two state variables. This creates four levels of confounding effects for this dataset, making it irreparably internally invalid. The first issue is that the parallel trends assumption has already been falsified to an extent, thus further assumptions to force results may be unwise. Secondly, to assess the difference in differences with group fixed effects, we must run two regression for both nodes due to the bidirectionally of the route. However, it is difficult to trust the results for a point estimate, since we are unsure how to decide which of the two estimates for the same outcome variable is "correct". Thirdly, because the data considers a directionless route, there is already an underestimation of any effect, since the tax effect only affect a proportion of the flights. Lastly, it is possible that by filtering the top 20 states, there are asymmetric effects for flights that both depart and arrive in a top 20 state and flights that only have one of the nodes within a top 20 state. While the first three are also present in the prior analysis, the final addition creates a layered issue which makes it non-meaningful to interpret further filtered results.

The relevant results can be seen in Table 7. We see changes in the average treatment effect in the average fare for all the panels. While they maintain their sign and magnitude, they all become insignificant at a 10% level. This is likely due to the lower statistical power of the DiD when considering fewer states leading to a larger standard error. We also see changes in the significance levels of the number of routes. The number of routes estimate for short haul routes stays significant at a 5% level, while maintaining the sign and a similar magnitude of the coefficient. The number of routes estimate for the long haul routes similarly is only now significant at a 10% level, however maintains it's sign and magnitude. This indicates the number of routes estimate is quite robust. Furthermore, there are changes to the coefficients in column 3. In Panel A we see a decrease in the statistical significance level from 1% to 10%. Additionally, we see quite a large decrease in the magnitude of the coefficient, despite it maintaining it's sign. In Panel C, we see a similar transformation, however here the coefficient is not statistically significant at a 10% level anymore. The average fare and the number of passengers can thus be considered as less robust of a smaller sample size adding uncertainty to the validity of the estimates. In contrast, the coefficients estimating the effect on the number of routes remain consistent and statistically significant.

Table 7: Difference in Differences Regression Results for 20 Most Populous States

Panel A: Total				
Variable	Average Fare (1)	Number of Routes (2)	Passengers Served (3)	ln (Passengers) (4)
ATET	-5.399	0.861	52009.969*	-0.021
	(4.438)	(1.631)	(27392.360)	(0.060)
cons	176.503***	65.422***	697882.823***	12.029***
	(5.475)	(1.125)	(34324.506)	(0.071)
Mean (Control)	202.215	67.480	1613463	12.356
N	15177	15177	15177	15177
Panel B: Short Haul				
Variable	Average Fare (1)	Number of Routes (2)	Passengers (3)	ln (Passengers) (4)
ATET	3.281	2.034**	2267.491	0.094
	(5.804)	(0.826)	(14386.664)	(0.086)
cons	150.427***	23.753***	286012.005***	11.163***
	(6.036)	(0.609)	(17973.744)	(0.074)
Mean (Control)	184.678	24.816	673495.1	11.381
N	14908	14908	14908	14908
Panel C: Long Haul				
Variable	Average Fare (1)	Number of Routes (2)	Passengers (3)	ln (Passengers) (4)
ATET	-5.536	-1.745*	41869.182	-0.254***
	(4.402)	(0.974)	(24234.715)	(0.055)
cons	204.624***	43.429***	430292.952***	11.397***
	(5.398)	(1.083)	(30206.073)	(0.083)
Mean (Control)	220.568	43.781	961800.2	11.810
N	14948	14948	14948	14948

Notes: Difference in difference regression estimation with the respective outcome variables for the 20 largest states omitting California. Time and Group fixed effects and interaction terms are included in each regression, however, the respective dummies are omitted for presentation purposes. The ATET is the coefficient of the treatment variable. The means of the control group are also provided for a better interpretation of the magnitude of the ATET by considering relative effects. Panel A considers total flights, Panel B considers only short haul flights and Panel C considers only Long Haul flights. Heteroskedasticity robust standard errors were used. "N" provides the sample size. Stars indicate significance where * $p \leq 0.1$, ** $p \leq 0.05$, *** $p \leq 0.01$.

8 Conclusion

This article discusses the effects of a decrease in the jet fuel tax in California on various supply outcomes. It shows rather strong and economically significant effects on the average fare, where a tax decrease reduces fares by up to 3%. This is complimented with an increase in the number of passengers. These results support a competitive market framework to describe the aviation market in the US. Within a standard supply-demand framework, a tax decrease on supplier shifts the supply curve to the right, validating both outcomes on fares and passengers. The magnitude of these findings suggest that a jet fuel tax is an effective policy instrument to affect consumer outcomes such as price and airline use.

When considering heterogeneity, some further nuance about policy outcomes come to light. Within the short haul market segment, airlines do not transmit the effect of the tax through price but rather through the number of routes. This is due to economic viability being more elastic to small changes in market conditions in contrast to longer-haul flights. Therefore, governments must consider geographic conditions and market prevalence of short haul flights to target it's polices for intended aims. A government which seeks to reduce prices in short-haul dominated market may have counterproductive outcomes by reducing jet-fuel taxes.

There are some limitations of this research. Firstly, and perhaps most importantly, the DiD method does not allow for statistical assessments of changes in trends post intervention. A key example of this was the ATET on the number of passengers when considering only long-haul flights, which misleadingly showed a negative coefficient, despite a short run positive but decaying effect. These nuances in effects may be better considered statistically with a panel regression or an event study. The second key limitation concerns the route-level data. Issues regarding both validating the common trends assumption, as well as the directionless city-pairs pose severe threats for the internal validity of the DiD estimates. A further study considering market shares, not only of the cheapest and largest airlines, but using competition indicators such as C4-index might provide a more comprehensive analysis on this statistic. A final limitation is the robustness to sensitivity analyses. When considering a subset of the total sample with lower statistical power, several coefficients became insignificant, plausibly alluding to the results being externally invalid in some instances.

Despite these limitations, this research is valuable in its contribution to policy effects on the supply side of the aviation market, a niche in literature that has been underdeveloped. Given a broader trend towards aviation based travel connecting various previously inaccessible parts of the world, governments are likely to have an enduring role in the market, encouraging their desired socio-economic outcomes. To that end, this paper demonstrates both the effects of a jet fuel tax, and the relevant considerations governments should consider when implementing it.

9 Bibliography

References

Administration, F. A. (2024). *Location identifiers search tool.* https://www.faa.gov/air_traffic/flight_info/aeronav/aero_data/loc_id_search/

Administration, U. E. I. (2024). *Jet fuel.* IndexMundi. https://www.indexmundi.com/commodities/?commodity=jet-fuel&months=240

Airlines for America. (2024). *A4a quarterly passenger airline cost index.* https://www.airlines.org/dataset/a4a-quarterly-passenger-airline-cost-index-u-s-passenger-airlines/

Bernardo, V., Fageda, X., & Teixidó, J. (2024). Flight ticket taxes in europe: Environmental and economic impact. *Transportation Research Part A: Policy and Practice*, *179*, 103892. https://doi.org/10.1016/j.tra.2023.103892

CDTFA. (2024). *Sales tax ates for fuel.* California Department of Tax and Fee Administration. https://www.cdtfa.ca.gov/taxes-and-fees/sales-tax-rates-for-fuels.htm

Cheung, T. K., Wong, C. W., & Zhang, A. (2020). The evolution of aviation network: Global airport connectivity index 2006–2016. *Transportation Research Part E: Logistics and Transportation Review*, *133*, 101826. https://doi.org/10.1016/j.tre.2019.101826

Chuang, S.-H. (2020). Aviation taxation and tax incidence. *Applied Economics*, *53*(4), 454–468. https://doi.org/10.1080/00036846.2020.1808179

Commission, E. (2021). *Taxes in the field of aviation and their impact* (Report). CE Delft. Retrieved June 11, 2024, from https://cedelft.eu/wp-content/uploads/sites/2/2021/03/CE_Delft_7M16_taxes_in_the_field_of_aviation_and_their_impact.pdf

Dama, A. A., Dequiedt, V., de Ubeda, A.-A., & Rota-Graziosi, G. (2023). *Taxation of civil aviation fuels as a source of financing for vulnerable countries* (FERDI Working Paper No. P318). Fondation pour les études et recherches sur le développement international (FERDI). Clermont-Ferrand. https://hdl.handle.net/10419/276971

Division, A. (2019). *Aviation in california fact sheet.* Contra Costa County. https://www.contracosta.ca.gov/DocumentCenter/View/60029/Aviation-in-California-Fact-Sheet--2019

Fukui, H., & Miyoshi, C. (2017). The impact of aviation fuel tax on fuel consumption and carbon emissions: The case of the us airline industry. *Transportation Research Part D: Transport and Environment*, *50*, 234–253. https://doi.org/10.1016/j.trd.2016.10.015

Gil, R., & Kim, M. (2021). Does competition increase quality? evidence from the us airline industry. *International Journal of Industrial Organization*, *77*, 102742. https://doi.org/10.1016/j.ijindorg.2021.102742

Gossling, S., & Upham, P. (Eds.). (2009). *Climate change and aviation: Issues, challenges and solutions* (1st). Routledge. https://doi.org/10.4324/9781849770774

Hubert, T., Guo, C., Mouton, C. A., & Powers, J. D. (2015). Tankering fuel on u.s. air force transport aircraft: An assessment of cost savings [Retrieved 21 November 2022]. *RAND Corporation.* https://www.rand.org/pubs/research_reports/RR837.html

Johnston, A., & Ozment, J. (2013). Economies of scale in the us airline industry. *Transportation Research Part E: Logistics and Transportation Review*, *51*, 95–108. https://doi.org/https://doi.org/10.1016/j.tre.2012.09.002

Karlsson, J., Odoni, A., & Yamanaka, S. (2004). The impact of infrastructure-related taxes and fees on domestic airline fares in the us. *Journal of Air Transport Management*, *10*(4), 283–291. https://doi.org/10.1016/j.jairtraman.2004.03.006

Keizer, R. (2024a). Consumer airfare report: Table 1 - top 1,000 contiguous state city-pair markets [Data Provided by Department of Transportation Office of the Assistant Secretary for Aviation and International Affairs. License: Public Domain U.S. Government. Accessed: 2024-06-22]. https:

//data.transportation.gov/Aviation/Consumer-Airfare-Report-Table-1-Top-1-000-Contiguo/4f3n-jbg2/about_data

Keizer, R. (2024b). Consumer airfare report: Table 7 - fare premium [Data Provided by Department of Transportation Office of the Assistant Secretary for Aviation and International Affairs. License: Public Domain U.S. Government. Accessed: 2024-06-22]. *Office of the Secretary of Transportation*. https://www.transportation.gov/policy/aviation-policy/competition-data-analysis/research-reports

Koopmans, C., & Lieshout, R. (2016). Airline cost changes: To what extent are they passed through to the passenger? *Journal of Air Transport Management*, *53*, 1–11. https://doi.org/10.1016/j.jairtraman.2015.12.013

Lijesen, M. G., Nijkamp, P., & Rietveld, P. (2002). Measuring competition in civil aviation. *Journal of Air Transport Management*, *8*(3), 189–197. https://doi.org/10.1016/S0969-6997(01)00048-5

MCICO. (2024). *Jet fuel vs. gasoline*. Maritime Capital International Corporation. https://www.mcico.com/resource-center/articles/jet-fuel-vs-gasoline

Ryley, T., Baumeister, S., & Coulter, L. (2020). Climate change influences on aviation: A literature review. *Transport Policy*, *92*, 55–64. https://doi.org/10.1016/j.tranpol.2020.04.010

Scowell, C. (2012, September). *A review of the aviation industry, 2008-2011 - dot oig*. Office of Inspector General. https://www.oig.dot.gov/sites/default/files/Aviation%20Industry%20Performance%5E9-24-12.pdf

Smyth, M., & Pearce, B. (2006). *Airline cost performance*. International Air Transport Association. https://web.archive.org/web/20150226013033/http://www.iata.org/whatwedo/Documents/economics/airline_cost_performance.pdf

Sobieralski, J. B., & Hubbard, S. M. (2020). The effect of jet fuel tax changes on air transport, employment, and the environment in the us. *Sustainability*, *12*(8), 3352. https://doi.org/10.3390/su12083352

White, Q., Agrawal, D. R., & Williams, J. W. (2019). Taxation in the aviation industry: Insights and challenges. *Transportation Research Record: Journal of the Transportation Research Board*, *2673*(9), 666–673. https://doi.org/10.1177/0361198119846102

Winchester, N., Wollersheim, C., Clewlow, R., Jost, N. C., Paltsev, S., Reilly, J. M., & Waitz, I. A. (2013). The impact of climate policy on us aviation. *Journal of Transport Economics and Policy*, *47*(1), 1–15. Retrieved June 8, 2024, from http://www.jstor.org/stable/24396349

Wuebbles, D., Gupta, M., & Ko, M. (2007). Evaluating the impacts of aviation on climate change. *Eos Trans. AGU*, *88*(14), 157–160. https://doi.org/10.1029/2007EO140001

Zachariah, R., Sharma, S., & Kumar, V. (2023). Systematic review of passenger demand forecasting in aviation industry. *Multimedia Tools and Applications*, 1–37. https://doi.org/10.1007/s11042-023-15552-1

10 Appendices

A California Maps

Figure 2: California on the US map

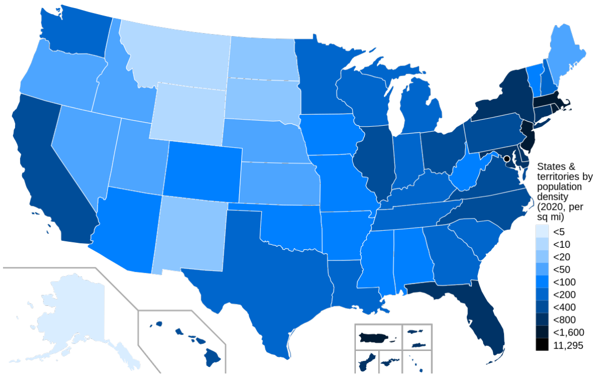

Figure 3: US States population density

B Parallel Trends

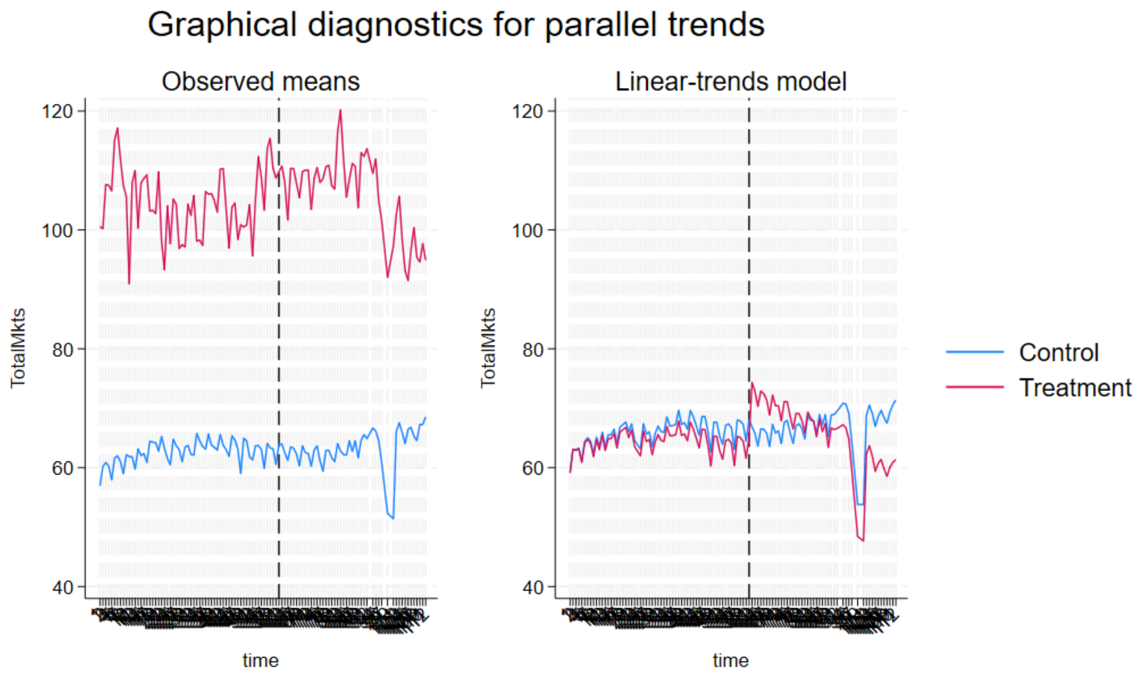

Figure 4: Common Trends for Total Routes

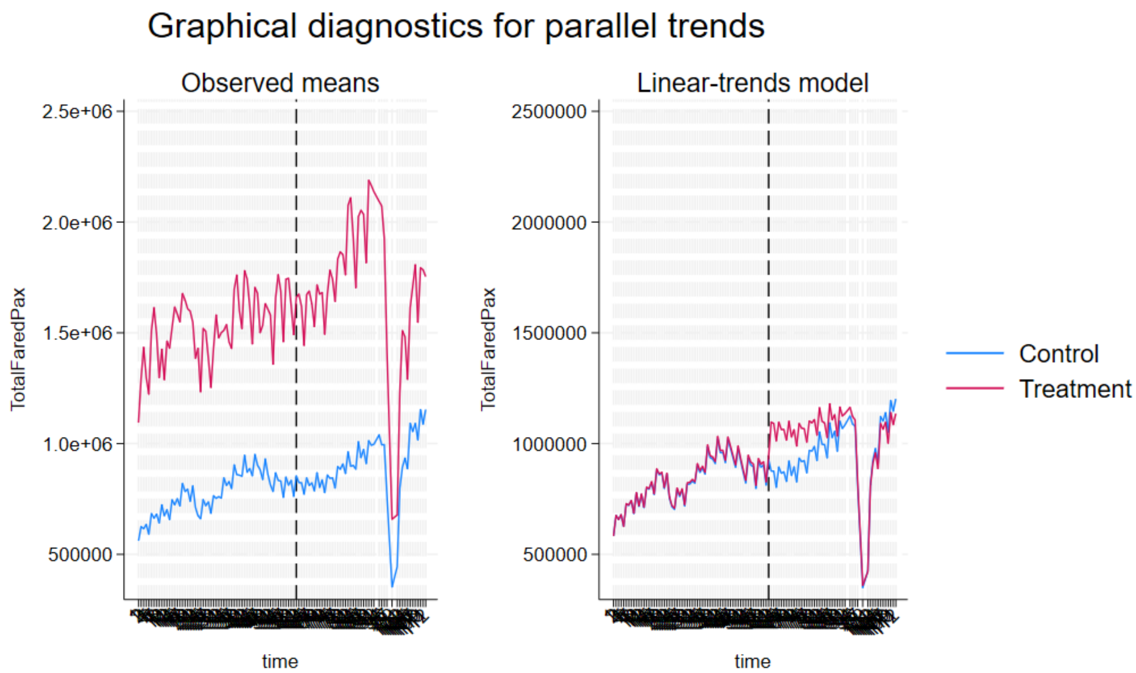

Figure 5: Common Trends for Total Passengers

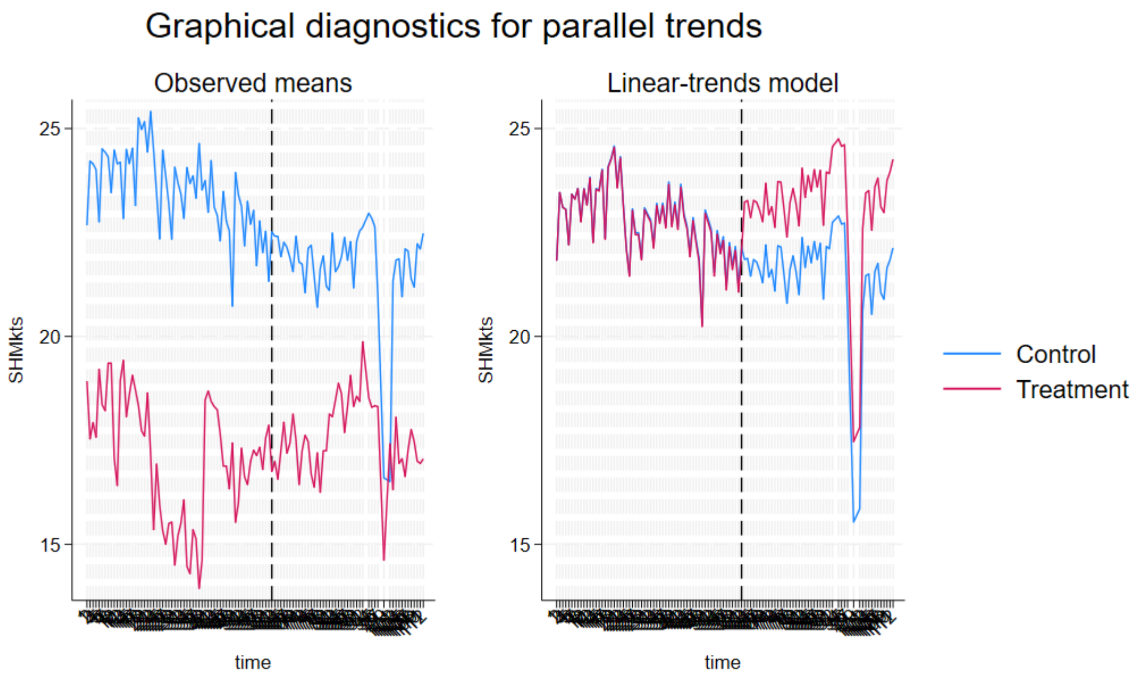

Figure 6: Common Trends for Short Haul Number of Routes

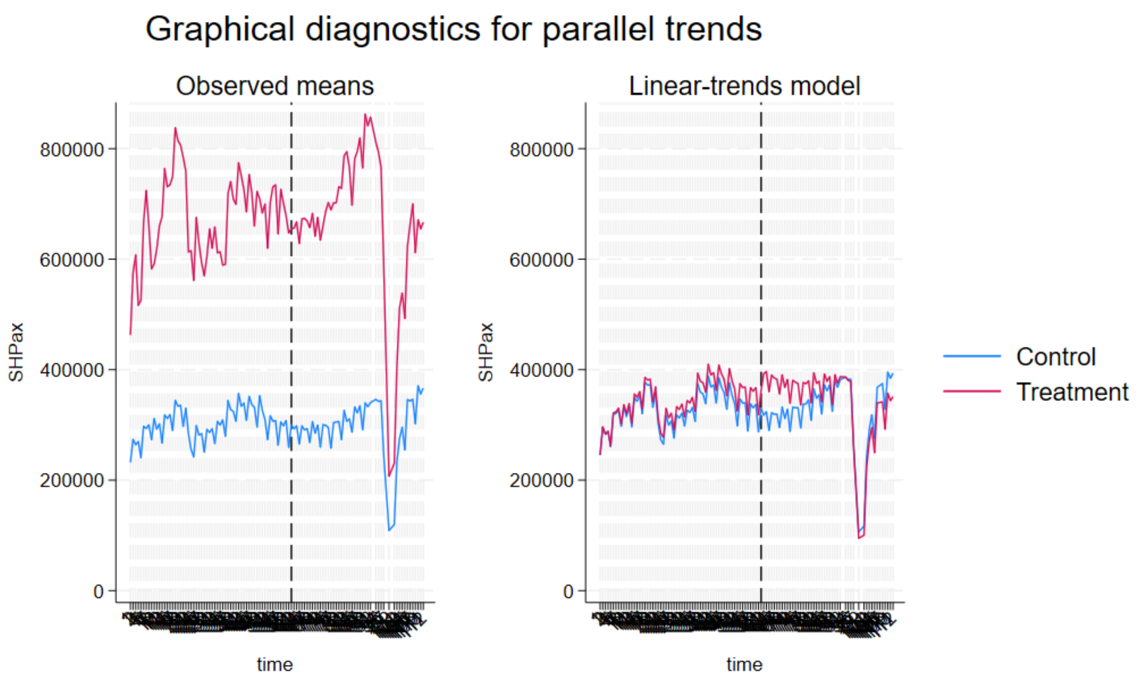

Figure 7: Common Trends for Short Haul Passengers

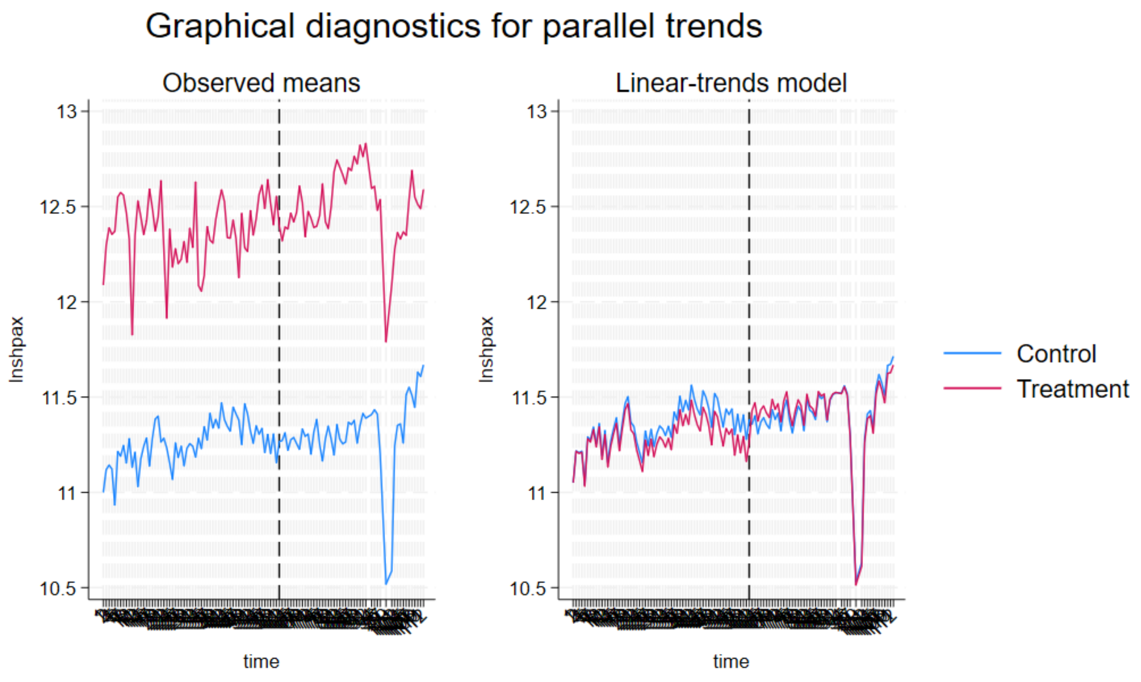

Figure 8: Common Trends for the Natural Logarithm of Short Haul Routes

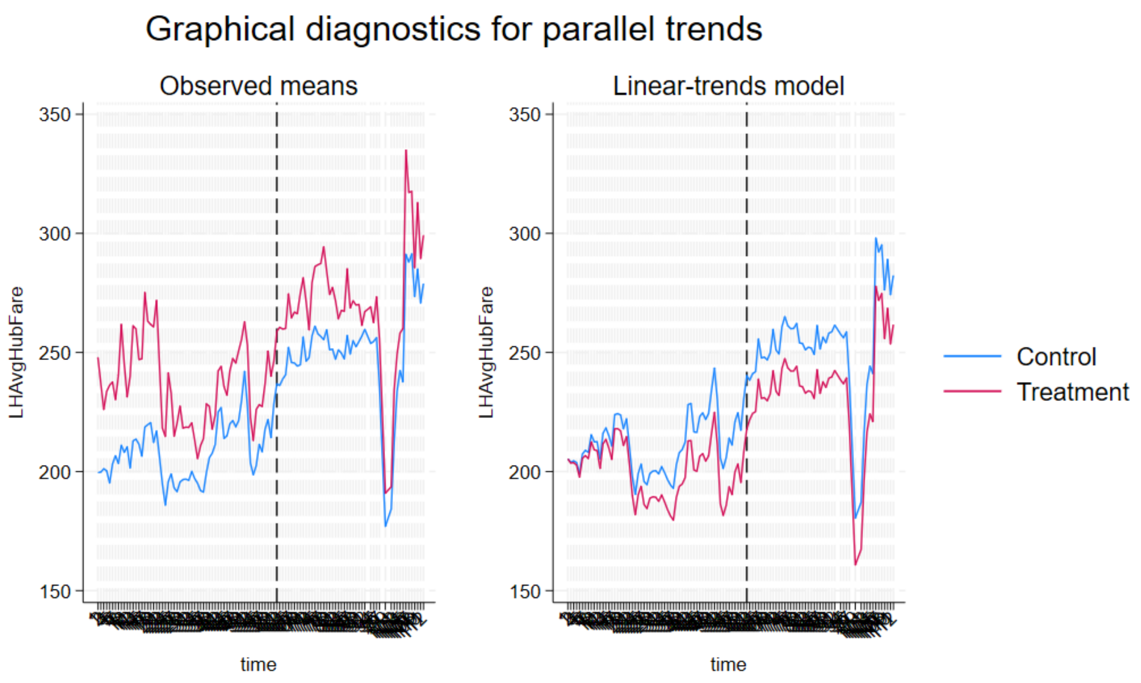

Figure 9: Common Trends for the Long Haul Fare

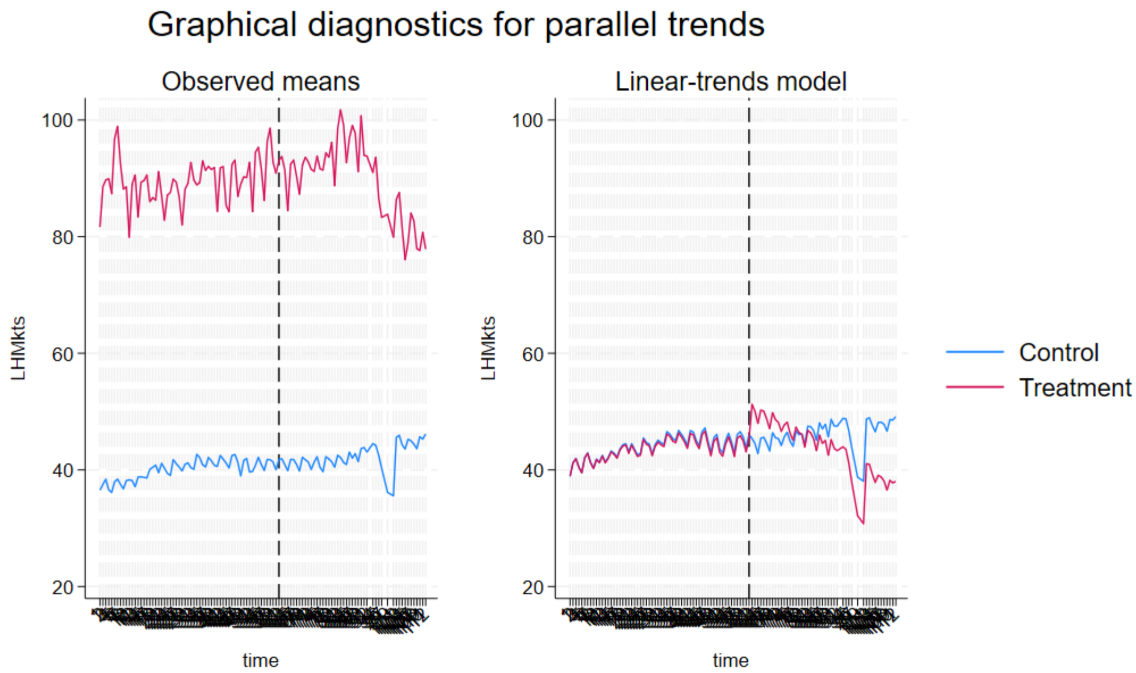

Figure 10: Common Trends for Long Haul Number of Routes

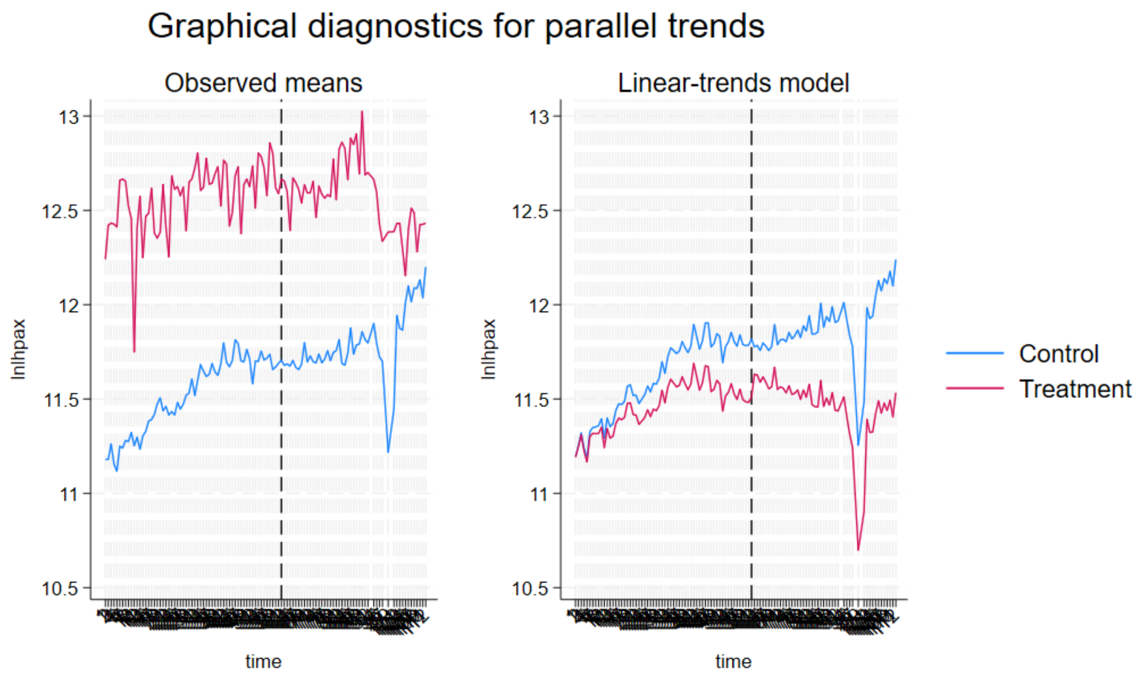

Figure 11: Common Trends for Long Haul Logarithm of Passengers