The Effects of the Swedish Aviation Tax on the Demand and Price of International Air Travel

(working paper)

Jonathan Stråle

Swedish University of Agricultural Sciences, SLU
Department of Economics
Working paper series, 2021:02
2021

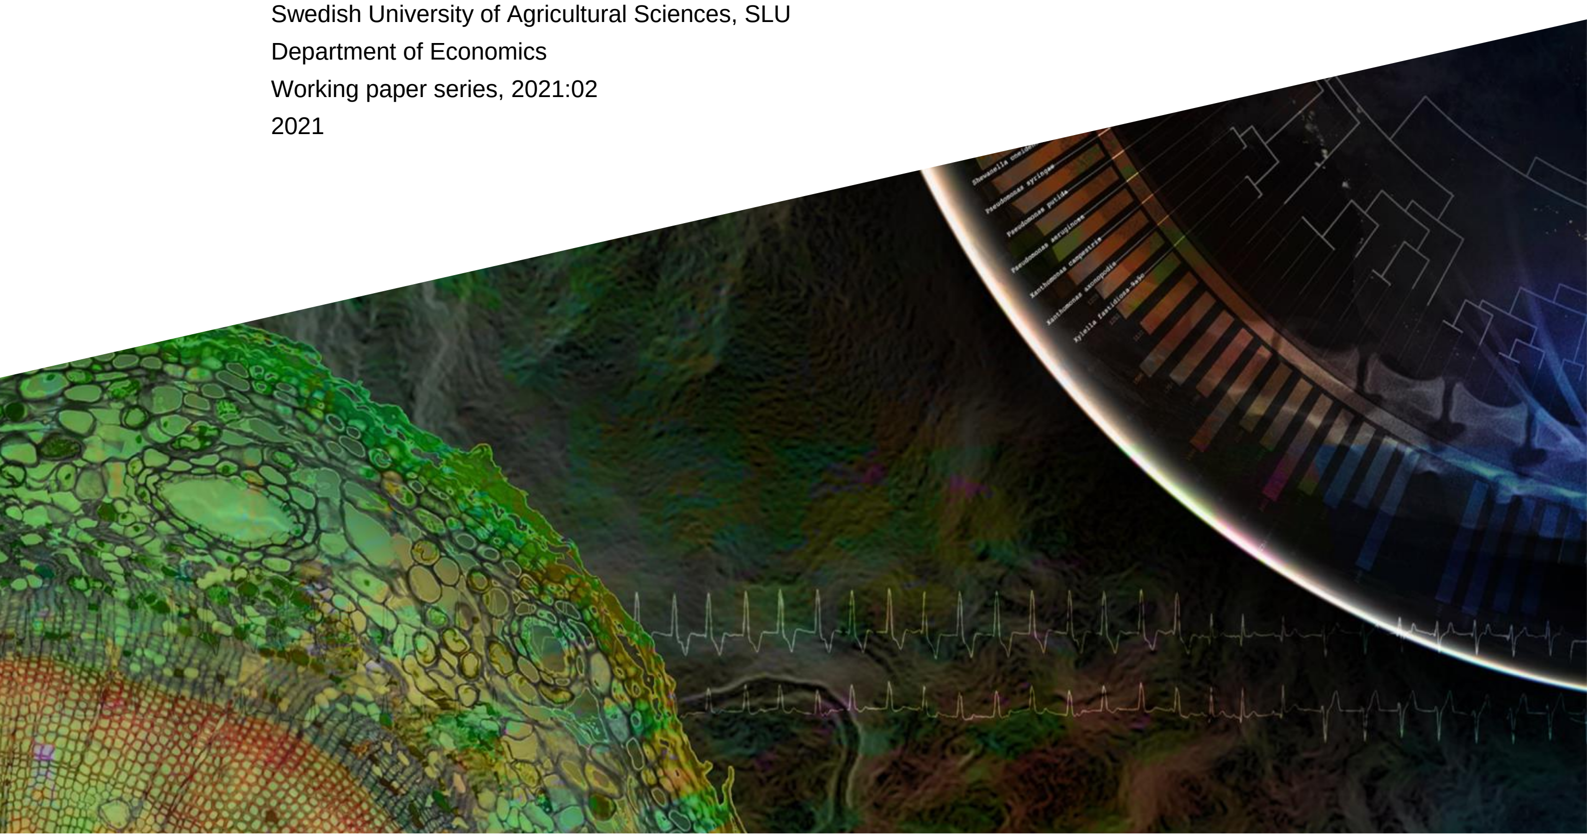

The Effects of the Swedish Aviation Tax on the Demand and Price of International Air Travel

Jonathan Stråle
Swedish University of Agricultural Sciences, Department of economics

Publisher:	Swedish University of Agricultural Sciences, Department of Economics
Year of publication:	2021
Place of publication:	Uppsala
Title of series:	Working paper series, Department of Economics
Part number:	2021:02
ISSN:	1401-4068
ISNR:	SLU-EKON-WPS-21/02-SE
Keywords:	Aviation tax, tax effects, synthetic control group method, price elasticity

Abstract

This paper evaluates the effects of the Swedish aviation tax on the price and demand for international air travel using the synthetic control group method. In addition, it estimates the price elasticity of international air travel from Sweden using web-scraped price data and an instrumental variable approach to account for the simultaneity bias in the price and passenger demand relationship. The effects of the tax are bigger than expected, with the effects on passengers increasing over time after the introduction and the effects on the prices starting high and then decreasing. The estimated price elasticity is -0.76, which together with the price effects of the tax accounts for the reduction in international travel the first three quarters after the introduction. The increase in passenger effects while the price effects diminishes indicates that there in addition to the price effect of the tax is a symbol effect that are affecting the behavior of the Swedish travelers. A potential "Greta Thunberg"-effect is also considered, and no direct evidence for such an effect are found. In contrast to other recent papers, no "leakage effect" of the tax, i.e. that passengers avoid the tax by going to neighboring countries, is found.

1 Introduction

Emissions from air travel continue to increase globally, mainly due to a strong positive trend in the number of passengers that travel by air. Sweden is one of the countries that travels the most by air, per capita, in the world, and in an attempt to hamper the continued increase in air travel demand the Swedish government introduced an aviation tax in April 2018. The tax is a fixed tax per passenger and the level of the tax is distance based, starting at around 6 Euros (60 SEK) for trips within Europe and ending at 40 Euros (400 SEK) for trips that are longer than 6000 km. The question is however, what is the effect of the tax on the demand for air travel? And what is the effect on the price of air travel? This paper tries to answer these questions through the use of a synthetic control group method (SCM), where Sweden is used as the treatment group and other European countries are used as potential donor countries for a synthetic control group. The effects on outbound international passenger numbers and prices are estimated. To do this, quarterly data from Eurostat on the number of passengers traveling internationally is used, together with a quarterly price index. In addition to this, the price elasticity of international air travels is estimated using web-scraped route level price data and an instrumental variable (IV) estimation where a cost shifting variable based on route distance and jet fuel price is used as an instrument for the price of air travel.

While several countries have introduced different types of air travel taxes, the economic literature of its effects is relatively scarce. The estimated effects in the literature are also somewhat conflicting. A few examples include Borbely (2019) and Falk and Hagsten (2019) who analyze the German and Austrian aviation tax. Borbely (2019), who employs the synthetic control group method at airport level, find that the tax, which is of a very similar size as the Swedish tax, reduces the total number of passengers from Germany with 2 percent, and that this effect can mostly be explained by the effect at low-cost carrier dominated airports as well as a leakage of passengers to airports close to the German border. Falk and Hagsten (2019) use a panel difference-in-difference approach to evaluate the effect of the tax in both Germany and Austria. They find bigger effects than Borbely (2019) as the effects they estimate are a reduction in international passengers by 9 percent the first year and 5 percent the second year after the introduction of the tax. They also find no evidence of passengers choosing airports in neighboring countries to avoid the tax, but they do also find that the main driver of the effect is found at low cost-carrier airports. Seetaram et al. (2014) uses an auto-regressive distributed lag model to estimate income, price and tax elasticities of route level international travelers from the UK, and find humble tax elasticities for most destinations, indicating a low effect of the tax on the number of passengers. There are also papers who use a modeling approach to simulate the effects of aviation taxes, using different price elasticities and assumptions of the price pass-through of the tax and substitutbility between modes of travel and domestic and international travel. Mayor and Tol (2007) do such an analysis where they evaluate different proposed changes to the air passenger duty, and find in general small effects of the tax on both passengers and emissions. Forsyth et al. (2014) is another example. They analyze the Australian aviation tax using a simulation modeling approach, especially focusing on how the tax affects the tourism industry in Australia, and finds that the industry will indeed suffer, the severity of which will depend on the actual price elasticity of air travel.

In contrast, price elasticities of air travel is relatively well researched. And there is good reason for this as price elasticities have been shown to vary a lot between e.g. source markets, destinations, type of travelers (e.g. business and leisure travelers) and the type of data as well as type of methodology that is used in the estimation (see e.g. Brons et al. (2002) and Peng et al. (2015)). While there is some consensus in the level of average price elasticities (in almost all cases the elasticity is negative, as expected), the width of the estimates is large, ranging from small positive elasticities in some cases to being well below a negative unit elasticity in other cases (Brons et al. (2002) show that elasticities vary between 0.21 and -3.2 in their meta study). When looking at the European market for international air travel however, the average elasticities are in most cases close to a negative unit elasticity. Most papers in the field, especially the older papers, also assume that price changes are exogenous, which given the pricing strategies of airlines is now known not to be the case. Prices of a certain flight are changed depending on the demand, which in itself depends on the price, and so forth. The dominating strategy to avoid this endogeneity bias in the literature is to find a valid instrument that correlates with the prices of air travel, but not with the choice to travel by the passengers. A recent example of this is Mumbower et al. (2014) who use web-scraped price and seating data to estimate price elasticities in the US market, and how it varies depending on a number of factors, while also addressing the aforementioned endogeneity problem using the air-lines prices in all other markets than the market that they study, as well as the number of daily nonstop flights that the competitor has, as instruments for the price. They find an average price elasticity that varies between -1.32 and -1.97 depending on the price, but that the elasticity also varies greatly on a number of other factors such as departure day of the week and time of day that the ticket is booked. Hsiao and Hansen (2011) is another example and they use the distance of the route times the price of jet fuel as an instrument, finding elasticities between -1.05 and -2.97 depending on route and market. Given this variety in estimates, it is important to estimate a trustworthy elasticity for the market that is being considered. For Sweden, the available literature is scarce. To the best of my knowledge, only one published paper exists, namely Kopsch (2012) who use aggregate price index data in their estimations and find price elasticities between -0.67 and -0.85

This paper contributes to both the literature on the evaluation of aviation taxes, a literature that still is rather sparse, as well as the literature on price elasticities air travel. Given that the evidence from the existing literature on the effects of aviation taxes show different effects, sometimes even contradicting effects, it is important to study the effects of these type of taxes further, not at least in other countries than has previously been researched. To the best of my knowledge, no paper have been published where the Swedish aviation tax has been evaluated. Sweden is also be an interesting case compared to e.g. Germany and Austria, as the substitutes to international air travel, such as trains, are not as available in Sweden compared to these countries. Understanding how effective aviation taxes is to hamper the demand of air travel is especially important given how emission-intensive air travel is, and as international air travel has been steadily increasing up until the Covid-19 pandemic. The results are therefore useful for policy makers that aim to reduce the emissions from the aviation sector. Similarly, as price elasticities differ depending on source country, and as only one other paper have been published where Swedish price elasticities are estimated, this paper adds to the literature on price elasticities of international air travel as well. Just as for the evaluation of the tax, this is useful for both policy makers and managers in the industry from both a Swedish and European perspective. By using web-scraped price data on route level while simultaneously handling the simultaneity bias in elasticity estimation we are also able to estimate a more detailed elasticity than most papers in the literature that rely solely on aggregate data, and in particular price indices instead of the actual offered price of a flight ticket.

The main findings in this paper are that the tax has a strong effect on the number of international passengers from Sweden, an effect that seems to be increasing over time. The effects are around 6 percent the first quarters and increasing to around 11 percent the last quarters in the time period studied (i.e. up until the start of the pandemic). The effects on the price of air travel is surprisingly high, given the size of the tax, the first quarters after the introduction of the tax (around a 10 percent increase), but then diminishes over time. The estimated price elasticity of air travel is -0.76 percent, suggesting that the increase in prices due to the tax can explain the drop in passengers in the first quarters after the introduction of the tax. Since the passenger effect is increasing over time as the price effect diminishes however, this indicates that there is also some sort of symbol effect of the tax that potentially increased the environmental awareness of the prospective passengers. While a potential "Greta Thunberg"-effect, which could be a possible explanation for the larger effect the second year, is investigated descriptively and no indication of a bias from her environmental messages can be found, this could potentially still be an explanation for the increased effect the second year after the tax. In the same manner, a relatively stronger environmental awareness with respect to air travel in Sweden compared to the control group countries could also be an explanation.

The rest of the paper is organized as follows. In the next section, the aviation tax is described. Section 3 presents the econometric strategy that is employed in the paper and section 4 describes the data that is used. Section 5 presents and discusses the results of the paper, section 6 provides a robustness analysis of the results and section 7 concludes.

2 The aviation tax

The tax is a fixed tax per passenger where the level is based on the distance of the final destination from Arlanda Airport. The tax is applied to all commercial flights, where the airplane is registered for carrying 10 passengers or more, taking off from a Swedish airport from April 1^{st} 2018 and onward. All passengers on the flights are included except infants (children younger than 2 years), transfer passengers from flights from outside of Sweden that are made within 24 hours, and the cabin crew. The responsibility for paying the tax is put on the airline, but the tax is usually communicated to the customers by the airline in the price specification before the booking is made. The tax-levels are, with a few exceptions, implemented as follows.

- Countries completely within Europe: 60 SEK / passenger (tax-zone 1)
- Countries completely or partially within another continent with a distance of at most 6000 km from Arlanda airport: 250 SEK / passenger. (tax-zone 2)
- Countries in a different continent than Europe with a distance of at least 6000 km: 400 SEK/passenger (tax-zone 3)

Some exceptions to this general rule exist as a result of a re-negotiation of the tax before it was voted through the Riksdag on November 22^{nd}, 2017. The detailed list of countries in the first and the second tax-zone respectively is found in 7, the countries in the third tax zone are the remaining countries.

3 Econometric strategy

To estimate the effect of the tax on the number of international passengers and the price of air travel, the synthetic control group method (SCM) developed by Abadie et al. (2010) and Abadie et al. (2015) is used. The main purpose of the SCM is to estimate a counterfactual outcome using a data driven approach that minimizes differences between the treatment group, in this case Sweden, and a control group that is constructed through a weighted average of countries that are similar to Sweden. By creating a control group that mimics the development of international air passengers, or the price index of air travel, before the tax is introduced in Sweden, counterfactual outcomes of international passengers and price index are estimated that represent the levels in Sweden if no tax had been introduced. The difference between Sweden and the synthetic control group after the tax thereby becomes the estimated effect of interest, under the assumption that the synthetic control group still successfully mimics the development in Sweden without the tax. Formally, following the terminology of Abadie et al. (2010) and Abadie et al. (2015), we have a sample of $J+1$ countries indexed by j, where unit $j=1$ is Sweden (the country of interest) and units $j=2$ to $j=J+1$ is the donor pool of potential control countries. There are $t=1,...,T$ time periods where time periods $t,...,T_0$ are the preintervention periods and $T_0+1,...,T$ are the postintervention periods. A synthetic control is represented by a $(J\times 1)$ vector of weights $W=(w_2+...+w_{J+1})'$, with $0\leq w_j\leq 1$ for $j=2,...,J$ and $w_2,...,w_{J+1}=1$. The choice of a particular value of W is thereby a choice of a synthetic control and the choice is made such that the characteristics of the synthetic control best replicates the treated unit's characteristics, where X_1 is a $(k\times 1$ vector of preintervention characteristics for the treated unit and X_0 is a $(k\times J)$ matrix of the same variables for all the countries in the donor pool. The optimal synthetic control W^* is thereby selected in order to minimize the difference between the preintervention characteristics of the treated unit and the synthetic control, represented by the vector X_1-X_0W. The variables that have the largest predictive power on international passengers and price index, respectively, are given the largest weights so as to best mimic the preintervention (and ultimately, the counterfactual postintervention outcome) development of these variables in Sweden. The estimated postintervention counterfactual outcome is thereby given by Y_CW^*, where Y_C is a matrix of the passengers or price index containing of the donor countries at a certain postintervention time t (borrowing the notation from Borbely (2019)). Letting Y_T be the outcome in the variable of interest in Sweden, international air passengers and the price index of air travel respectively, the estimated effect of the tax thereby becomes $Y_T-Y_CW^*$, for the time period after the implementation of the tax ($t > T_0$).

The recommendation by Abadie et al. (2010) and Abadie et al. (2015) is to only include countries in the donor pool where air travel as well as other predictors are not too dissimilar to the unit of interest, i.e. Sweden. Due to this, only countries within the EU are used as potential controls, and in addition, countries that have more than twice the number of international passengers are excluded from the donor pool of potential controls. In addition, countries that have had a similar treatment as Sweden during the time period of consideration needs to be excluded, as one of the identifying assumptions in the SCM is that the potential controls have not been exposed to a similar intervention as in the country of interest. While several countries already had aviation taxes in the beginning of the observed time period, only a few had taxes introduced or changed during the time period of study, namely: Norway and the UK. UK is already excluded because of the size of international air travel, but Norway is also excluded for this reason. Austria also has a tax that is introduced in 2011, but as this is in the very beginning of the time series considered (as is described in more detail in the next section), and we thereby only observe the post-tax trends, it is left in the donor pool.

As international travel is highly seasonal, both the passenger and price index variable is deasonalized before the SCM is used. This is done by estimating the average quarterly effect using a linear regression where quarterly dummy variables are used to estimate the average seasonal effect for the pre-treatment time. As the raw-data comparisons suggest that there might be an anticipation effect the summer before the tax is introduced, the desasonalization is based on the time period before the second quarter in 2017. The estimated average quarterly effect is then removed for the whole time series. This is done for all countries that are included, Sweden as well as all potential donor countries.

One drawback with the SCM compared to standard regression methods is that a confidence interval cannot be easily estimated, and thus inference is not as straight forward. To assess how likely it is that the estimated effects are due to a random fluctuation, rather than being a true effect of the tax, placebo estimations are made for the whole donor pool of potential control countries. The effect in Sweden is then compared to the "effects" that are estimated for these countries, where we expect the Swedish effect to be the most extreme (as no similar intervention has been made in the donor countries). A p-value is calculated by comparing the "extremeness" of the Swedish effect to these other countries, while at the same time adjusting for how well the pre-fit is made in these placebo estimations (as a poor pre-fit by itself indicates that any observed "effect" after the placebo-intervention time is less believable). This procedure yields what is called a standardized p-value, and should be interpreted as the probability for the observed effect to be a random event, after adjusting for the quality of the pre-fit in the placebo estimations, i.e. in a very similar fashion as the commonly used p-value in standard regression methods. As the calculation of the p-value is, in essence, a comparison of "extremeness" between the effects in Sweden and the placebo-effects in the donor pool, an adjusted p-value of 0 is possible. This does however not mean that it is a zero chance of the effect to be a random event, only that Sweden has the most extreme effect in comparison to the donor pool in that specific post-tax time period.

To estimate the the price elasticity of international air travel, an instrumental variable estimation is instead used. This is done in order to correct for the likely bias that stems from the potential reverse causality in the relationship between air travel ticket pricing and demand of air travel, i.e. that the price affects the demand for air travel at the same time as the demand for air travel also affects how airlines price their tickets. In order for the IV-estimation to solve this issue, the exclusion restriction must hold and the instrument must be relevant. The exclusion restriction requires the instrument to have no correlation with the dependent variable outside the correlation that goes through the explanatory variable of interest, the variable that is instrumented. Ideally, one wishes the instrument to be randomly assigned, but when observational data is used this is rarely the case. Instead one needs to find a variable that likely fulfills the exclusion restriction anyway, and as this is an untestable assumption the feasibility of this assumption must be examined by reason. That the instrument is relevant means that the instrument must have a significant effect on the variable that is instrumented, something which is easily tested in the first stage estimation. When it comes to possible instruments that can be used for the price of air travel, the literature proposes a few alternatives. Mumbower et al. (2014) provide a comprehensive discussion of the most commonly used instruments, and divide them into four different categories. One category of instruments is cost-shifting variables, a variable that influences the cost per trip for the airline but that does not affect the passengers decision to travel (outside the effect the variable has on the price). This variable is most commonly used for estimations using aggregate data, and some examples in the literature are Hsiao and Hansen (2011) who use the jet-fuel price times the route distance as an instrument (which is also used in this paper), Berry and Jia (2010) and Granados et al. (2012) who use a hub-indicator for the destination airports (as flights to hubs can be done at a lower marginal cost due to the use of more fuel efficient air carriers) and Granados et al. (2012) who use distance of the route as an instrument. Another category is referred to as Hausman-type instruments, where the price of the same brand or product in other markets are used as an instrument. An example being Gayle (2004) and Mumbower et al. (2014) who use the average price in all other routes that have a similar length as the route being studied. A third category is Stern-type instruments, where measures of competition and market power are used as an instrument for the price. Berry and Jia (2010) use the number of carriers offering flights in a certain route as an instrument for example and Granados et al. (2012) and Mumbower et al. (2014) use how many daily nonstop flights in the market being studies that are offered by competitors as an instrument. The fourth category presented by Mumbower et al. (2014) is measures of non-price characteristics of other products supplied by firms, the same firm as being studied or others, in the same market. An example being Berry and Jia (2010) who use the percentage of rival routes that offer direct flights, the average distance of rival routes and the number of rival routes as instruments for the price.

The one that is chosen in this paper is a cost shifting instrument that is proposed by Hsiao and Hansen (2011), namely the length of the route times the jet-fuel price in the month that is being considered. As the price of jet-fuel affects the costs for airlines, and the total cost per trip for the airline depends on the distance of the flight, it is likely to be correlated with the price of the ticket for that flight. Given that this is the case, which of course is easily tested, the instrument would be relevant (this is also the case in this paper). In addition, the instrument must fulfill the exclusion restriction, i.e. that it only affects the number of passengers on a flight through the price, and thereby is not correlated with the number of passengers in another way, when other covariates are controlled for. When it comes to the price of jet fuel, this makes sense. It is hard to see how the price of jet fuel would influence individuals decision to travel abroad. Potentially the jet fuel price could be correlated with the economic cycle in general, which of course is correlated with how many people travel abroad, but this can be controlled for through e.g. controlling for the employment rate or GDP/capita. When it comes to the distance of the journey, it is not as clear as the length of the journey is strongly correlated with the time it takes to reach the destination, which incurs a time cost on the passenger that acts as a damper of the demand for the trip. On the opposite side, destinations further are in general considered to be more exclusive and can act as a enhancer of the demand. In any case, the distance is fixed and should therefore affect the variation in demand for specific routes over the time period that is studied. All in all therefore, it can be argued that this instrument fulfills the exclusion restriction and that the instrument is valid.

The IV-estimation is conducted in two stages.. The first stage is given by

$$\hat{Z_{gt}} = \rho_{gt} + \gamma I_{gt} + \mathbf{X}_{gt} + \eta_{gt} \tag{1}$$

where $\hat{Z_{gt}}$ is the log of the ticket price per route and month (the monthly average of the daily price data), I_{gt} is the instrumental variable (the jet-fuel price the given month times the distance of the route), and $\mathbf{X}_{gt}$ is a vector of control variables including house price index, employment rate, GDP/capita, temperature in Sweden, and a Google search index variable for the term “Greta Thunberg”. Time fixed effects are also included. The second stage is given by

$$Y_{gt} = \theta_g + \delta_t + \beta\hat{Z_{gt}} + \mathbf{X}_{gt} + u_{gt} \tag{2}$$

where Y_{gt} is the log of the number of passengers per route and month, $\hat{Z_{gt}}$ is the projected price variable from the first stage and $\mathbf{X}_{gt}$ is the same vector of control variables that was used in the first stage. As both ticket price and number of passengers are log transformed, β can be directly interpreted as the average price elasticity for international air travel from Sweden, for the routes that are considered, given that the identifying assumptions hold.

4 Data

The main outcome variable is the number of international passengers from Sweden. The data is quarterly data and for the main analysis the aggregate passenger data on international air traveling for different countries will be used. In addition to number of passengers, data on how the prices of air traveling have developed is also used to examine a potential effect of the tax on the price of air travel. It consists of country level price indices where the most weight is put on the most popular routes, so as to most accurately portray the true average price of air travel. This price index captures all types of air travel, so not just international travel. The predictor variables that are used are GDP/capita, house price index, employment rate and the size of the population in the countries. All of these data sets are publicly available from Eurostat.

The data period that is considered in the main specification for the effect on passenger numbers starts in the first quarter of 2011, and ends in the last quarter of 2019. The start of the period is chosen as most of the idiosyncratic shock in international travel that resulted from the 2008 financial crisis, which affected different countries to a varying degree, has been recovered by 2011. Given that quarterly data is used, this also gives plenty of pre-intervention time periods for the SCM to work well, as well as doing tests of inference (such as in-time placebo checks). The end of the period is chosen as to avoid any bias from the extreme decline in air travel that inevitably resulted from the covid-19 pandemic. For the effect on the price index, a shorter time period, 2015–2019, is used due to issues with finding an acceptable pre-fit when longer time periods are used.

For the estimation of the price elasticity of air travel, route-level price data has been web-scraped from Expedia.com on a daily basis between January 2018 and January 2019. In order to get a representative price for the destination at a certain date, prices for the same take-off date was scraped 1 week, 5 weeks and 9 weeks before take-off (except for the first and last weeks of the web-scraping). One way tickets from Stockholm (Arlanda airport) to the main airports in Barcelona, Paris, Brussels, Luxembourg, Gran Canaria, Funchal, Kap Verde, Hurghada, New York, Hong-Kong and Phuket was scraped [1] In addition to price, the number of stops, total time of the flight and the name of the airline

[1] In addition, ticket prices for the same destinations but from Helsinki, Copenhagen, Oslo and London was also scraped but not used at this point in this paper.

was scraped. Due to data quality issues, only six of these routes could be used in the estimation, namely Paris, Barcelona, Bryssel, Luxembourg, New York and Hong Kong.

In addition to the predictor variables that are used in the synthetic control group method, which are used as control variables in the IV-estimation, some other variables are also included in the IV-estimation. The instrumental variable that is used in the price elasticity estimation is monthly jet-fuel price data multiplied by the route distance. The jet-fuel price data is taken from the US Energy Information Agency (EIA) and the route distances is taken from Google maps. Additional control variables are also included. The temperature in Sweden is included as a control variable as this variable possibly can affect both the pricing and demand of air travel, and the data is taken from the European Climate Assessment & Dataset project. In addition, a proxy variable for the break through of Greta Thunberg, whose environmental message grew in popularity towards the end of 2018, is used, which is Google trends data for searches on "Greta Thunberg" in Sweden.

5 Main results

One of the advantages with using the synthetic control group is that it is very easy to get a visual understanding of the estimated effects. In Figure 1 (a) and (c), the development of the number of international passengers and price index from Sweden, after de-seasonalization, is compared with the synthetic counterpart before and after the tax is implemented (the time period just before the introduction of the tax is indicated with a vertical line in the graphs). This is the main specification, where the pre-fit does not include the summer before the tax, as a potential anticipation effect is present in the raw data, and both Denmark and Finland is part of the donor group.

As can be seen, the synthetic control groups do a good job in mimicking the development of international passengers and the price index of air travel before the tax is implemented, and as soon as the tax is implemented there are clear differences between the development in Sweden and the synthetic control group, for both the number of international passengers and the price index. Figure 1 (b) and (c) show the differences between the development in Sweden and the development in the respective synthetic control group, a difference that can be read as an evaluation of how well the pre-fit of the synthetic control group is made for the time before the tax is implemented and how big the effects of the tax for the time period after the tax is implemented. The estimated effects corresponding to the graph are also presented in Table 1, the main specifications, together with the standardized p-value estimation, which can be interpreted as the probability of the result happening by pure chance. In the main specification, the tax seems to have an initial effect of almost 200 000 fewer passengers passengers the first quarter (or a reduction with 4.4 percent compared to the synthetic control group). The effect then seems to be

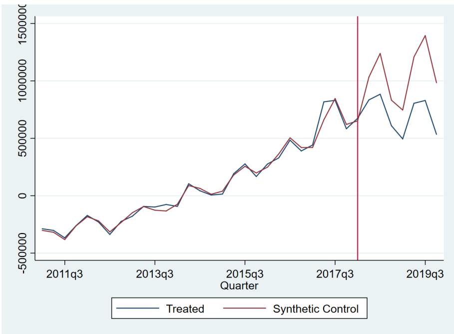

(a) Development of de-seasonalized international passengers in Sweden and synthetic control group

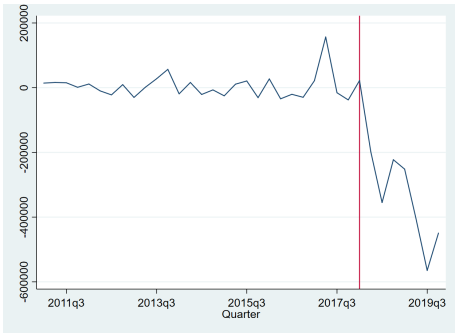

(b) Effects of tax on deseasonalized passengers

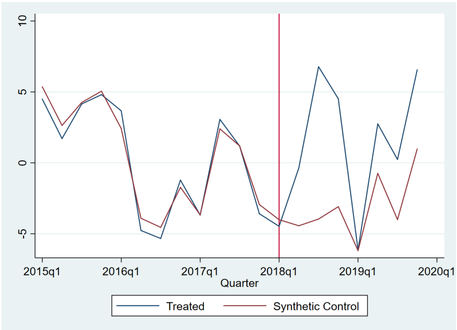

(c) Developement of deseasonalized air travel prices in Sweden and synthetic control group

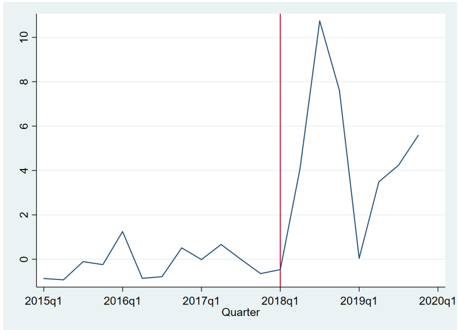

(d) Effects of tax on deseasonalized air travel price index

Figure 1: Effects of the air travel tax on international passengers and price index, de-seasonalized time series

increasing with time, being the largest in the third quarter the year after where the effect is around 560 000 fewer passengers (a reduction with about 11 percent compared to the synthetic control group). The effect on the air price index is sizable at first, being as high as almost 11 index points (10.4 percent) the second quarter after the tax is implemented, but after the first three quarters the effect is no longer significant (except for the last quarter in 2019). Apart from the first quarter in 2019, where the effect seems to be close to zero, the estimated effect for the remaining quarters is however positive, all though smaller than the initial effects and only significant in the last quarter of the year. Given the size of the tax, which is only about 6 Euros for travel within Europe (which is where most air travel is made from Sweden), these effects are surprisingly large. Especially the initial large effects on the price index of air travel is much larger than expected, given that a representative ticket to destinations within Europe costs between 130–270 Euro[2]. One potential explanation is that the tax gave the airlines operating in Sweden an opportunity for increased tacit co-operation on prices, and increased them with more than the amount of the tax as this is beneficial for all of them given that a price war can be avoided, and

[2]Based on the web-scraped data used in this paper, for the destinations within Europe.

Table 1: Quarterly effects of the Swedish aviation tax on passengers and price index of international air travel. Standardized P-values below estimated coefficient.

	Passengers (main)	Passengers (full fit)	Passengers (No FI or DK)	Prices (main)	Prices (no 2017q2)	Prices (No FI or DK)
2018q2	-198175.8	-336091.5	-280979.8	4.077296	4.970327	4.564919
	.047619	0	.0526316	0	.0952381	0
2018q3	-355040.8	-440988.8	-563527.3	10.73874	8.388437	10.624
	0	0	0	0	.0952381	0
2018q4	-222591.6	-453004.8	-321622.8	7.609439	2.858734	8.189187
	.047619	0	.0526316	0	.5238095	0
2019q1	-251717.8	-509518.6	-336985.8	.044559	-1.856733	.41039
	.047619	0	.0526316	1	.5238095	.6842105
2019q2	-403750.1	-693316.8	-507440.7	3.488055	.15097	4.259926
	.047619	0	.1052632	.1428571	1	.1052632
2019q3	-564845.3	-771135.9	-734069.8	4.23904	1.541584	4.374238
	0	0	0	.2380952	.8095238	.2631579
2019q4	-449329.3	-758368.5	-579018.2	5.582713	1.394078	5.356411
	0	0	0	.047619	.7619048	0

that this tacit co-operation could only last for a few quarters before the pricing strategies went back to normal.

To be able to evaluate whether the effect on passengers is reasonable given the price effects, we first need to estimate the price elasticity of international air travel. These estimates are found in Table 2, where the price elasticity is first estimated using a basic OLS-estimation, which is included as a base line comparison even though it is most likely biased, and then with the IV-estimation, which is the main price elasticity estimate. As can be seen, the estimated price elasticity from the IV-estimation is -0.758. This means that if the price of air travel increases with 1 percent, the demand for air travel will decrease with on average 0.758 percent, on the routes that have been studied. This is slightly lower than the average estimates in the literature, but as mentioned in the introduction the space of estimated price elasticities in the literature is very large and differs between countries.

Given the estimated price elasticity, we can also get an indication of how much of the tax effect that is a pure price effect and how much is due to a symbol effect of the tax, which is another channel that the tax can have an effect through. For the first three quarters we see a tax effect on the price index of an increase of 4.3, 10.4 and 8.5 percent which is accompanied by a reduction in passengers by -4.4, -7.2 and -5.9 percent. Given the estimated price elasticity and the effect of the tax on the price index, the decrease in passengers are very close to what is expected. For the second year, this is however not the case, as the price effects are smaller and not significant for most quarters at the same time as the passenger effects are larger than in the first year. This therefore implies that there are more than price effects of the tax that are at play for these quarters. A symbol effect of the tax is one potential explanation for the increased effect of the tax on passengers the second year after the tax implementation. As the tax was launched as an environmental tax, at the same time as it got a lot of media attention both before and after its introduction, it is possible that this attention made, at least some, Swedish travelers more aware of the environmental damage that air travel has and chose to travel less because of this. Given that the tax is constant over the period studied, at the same time as the effects increase, this seems plausible even without the price estimates. It could however also be that the tax is a result of an already increased environmental awareness, and that we would have seen a decrease in air travel anyway. Given the timing of the reduction in air travel, together with the higher than normal air travel the summer before the introduction of the tax (discussed further below), this is however not as likely. Finally, there is another potential explanation for the increased effect the second year after the introduction of the tax: nameley Greta Thunberg. As Greta rose to fame with her pro-environment message at the end of 2018 and beginning of 2019, it is possible that part of the relative reduction in air travel in 2019 is due to an increased awareness because of her. This is however considered further in section 6, where no direct evidence of a "Greta-effect" is found. Nevertheless, Swedes are in general a climate aware population and the environmental awareness have likely increased over the past years. Air travel has also received a lot of focus in the environmental debate, where words such as *flygskam* (flight shame) and "stay-on-the-ground" campaigns in social media very likely have affected the demand for air travel. It is therefore not unlikely that the bigger effect the second year can be in part explained by this increase in environmental awareness, which might have been relatively larger than in the countries in the control group.

As the calculation of the p-values that are used for the inference for the tax effect estimates are a bit different from ordinary regression methods, it is valuable to discuss these a bit more in depth. These p-values are based on a series of placebo studies where each potential donor country are considered as the treatment country, instead of Sweden, and placebo effects are calculated. The actual effect when Sweden is the treatment country is then compared to these other placebo effects, where one expects the true treatment effect to be an "outlier" in the sample of estimated effects. The results of these placebo studies are given in Figure 2, where the difference between the treated and synthetic control group is presented for all the potential donor countries and Sweden. As can be seen the actual effect in Sweden is part of the most "extreme" effects in both the passenger and price estimations, which is what is needed for significant effects. However, the placebo

Table 2: The price elasticity of Swedish international air travel

VARIABLES	First stage	OLS-estimation	IV-estimation
Ticket price, log-form		-0.620***	-0.758***
		(0.127)	(0.130)
Employment rate	0.871	0.154	0.197
	(0.825)	(0.387)	(0.352)
House price index	-0.153	-0.0414	-0.0651
	(0.117)	(0.206)	(0.186)
Greta effect	0.0371	0.00620	0.0102
	(0.0275)	(0.0639)	(0.0577)
Temperature	-0.00569	-0.000933	-0.00166
	(0.00526)	(0.0115)	(0.0103)
IV	0.785***		
	(0.0360)		
Constant	-49.35	5.678	6.005
	(50.54)	(19.09)	(17.20)
Seasonal dummys	YES	YES	YES
Observations	75	67	67
R-squared	0.872	0.335	0.321

Robust standard errors within parentheses

*** p<0.01, ** p<0.05, * p<0.1

estimation for passengers looks far more neat than the equivalent for prices: in particular, the pre-treatment fit for the passenger estimation is much more exact than for the price estimation, where the donor pre-treatment fits are heavily scattered. In the calculated p-value this is taken into account as the calculated probability is standardized by the how well the pre-fit in the placebo estimations are. The estimated p-value is thereby an estimate of the probability that the estimated effect in Sweden happened purely by chance, where a zero indicates that there is not other country that had as big of an effect when the pre-treatment fit is accounted for. Hence, even though there seems to be a few countries that have more "extreme" effects in Figure 2, when the pre-fit is accounted for the Swedish effect is most "extreme", or outnumbered only by one of the placebo treatments, for the quarters that show significant results in Table 1. Even though the p-values correct for how well the pre-fit is made in the placebo estimations, the volatility of the development in price indices in all countries warrants some caution in the interpretation of the estimated effects of the tax on the price index in Sweden as the underlying volatility in the price indices increases the risk of the observed effect being a random event. In addition, a lot of pre-treatment predictor variables of the actual outcome variable, the price index, are needed to achieve an acceptable pre-treatment fit, as can be seen in Table 4, which increases the risk of over-fitting. This is not surprising as the price of air travel is notoriously volatile, but it does also warrant some caution in the interpretation of the the tax effect on the price level of air travel, both in terms of size and significance.

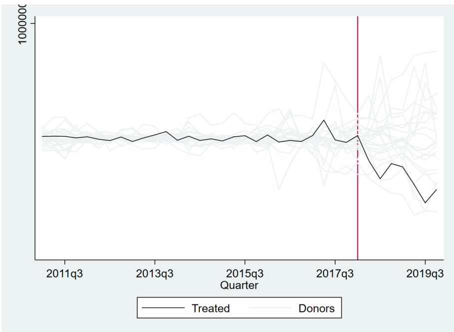

(a) Placebo estimations of deseasonalized international passengers

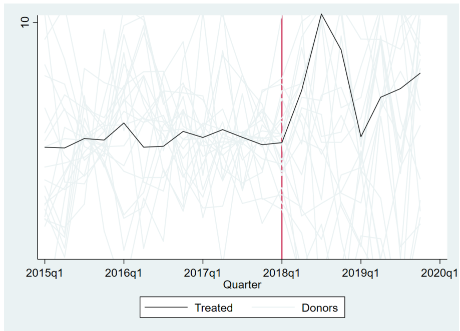

(b) Placebo estimations of deseasonalized price index

Figure 2: Placebo estimations for other countries to calculate the standardized p-value.

When interpreting these results, it is also important to take into account how well the variables that are used as predictors are balanced between the treatment group, Sweden, and the synthetic control group. Ideally we would like to see a perfect balance, that all used predictor variables are of the same value in the treated and synthetic group, as this would provide further confidence to the interpretation that the synthetic control group is comparable to Sweden and that the passenger numbers and prices would have developed in the exact same way as the control group, had not the Swedish tax been implemented. As can be seen in Table 3 and 4, this is unfortunately not the case for the predictor variables that are not the included lags of the variable of interest (passengers or price index). Average quarterly GDP/capita and employment rate in the used pre-treatment period is for example higher in Sweden in both the passenger and price estimation. The balance for the included lags of the main variable of interest are in both cases quite good however, and it is therefore clear that these are the most important predictors to achieve the fit that we see in Figure 1. While not ideal, an imperfect predictor variable balance is only an issue if the variables are really important for the development of the main variable of interest (passengers and/or price) as well as having a different development in the post-treatment period compared to the treatment group. If the development of the variables are the same in the synthetic control group and Sweden, both before and after the tax is introduced, then this unbalance in the predictor variables is no cause for concern. This has been considered, and for all the predictor variables the trend is very similar to Sweden both before and after the introduction of the tax, as can be seen in Appendix A.3. Hence, the observed imbalance in the predictor variables should not introduce any significant bias to the estimated effects.

Table 3: Balance of predictor variables, passenger estimation. Pre-tax averages is only up until 2017q2.

	Treated	Synthetic
De-seasonalized passengers 2012q1	-171900.8	-183293
De-seasonalized passengers 2012q3	-337606.8	-315415.5
De-seasonalized passengers 2013q3	-98716.75	-126266.7
De-seasonalized passengers 2016q3	276991.3	256209.9
De-seasonalized passengers 2016q4	389663.8	419328
House price index, pre-tax average	89.3812	99.18674
GDP/capita, pre-tax average	11347.23	7704.241
Population, pre-tax average	9653101	1.19e+07
Employment rate, pre-tax average	74.76	65.41955
Ticket price index, pre-tax average	98.914	101.9792

Table 4: Balance of predictor variables, price estimation. Pre-tax averages is only up until 2017q2.

	Treated	Synthetic
De-seasonalized price index, 2015q3	4.151109	4.258269
De-seasonalized price index, 2015q4	4.806666	5.049587
De-seasonalized price index, 2016q3	-5.335557	-4.549818
De-seasonalized price index, 2017q2	3.063336	2.398287
De-seasonalized price index, 2017q1	-3.683336	-3.668129
De-seasonalized price index, 2017q3	1.184443	1.188242
De-seasonalized price index, 2017q4	-3.59	-2.943606
De-seasonalized price index, 2018q1	-4.47667	-4.008589
Total air passengers, pre-tax average	5389587	1179040
House price index, pre-tax average	108.3015	108.5405
GDP/capita, pre-tax average	11773.35	8147.883
Employment rate, pre-tax average	76.19231	73.83655

Turning back to Figure 1 (a) and (b), there is one particularly interesting observation to be made in the pre-treatment time period: namely that there seems to be a positive effect in the second quarter the year before the tax was introduced. As this "anomaly" was observed already in the raw-data for Sweden, it is no surprise that it is still present in the synthetic comparison as the pre-treatment fit was done on the value of the used predictor variables before this quarter (so as to not inflate the tax effect artificially by forcing a fit based on this uncharacteristic large number of passengers the summer before the tax, as described in 3). This is especially interesting as this is the quarter when the tax was announced, and it might thus be a behavioral effect of either the passengers or the airlines that we are observing. In both cases the potential behavioral response could be that it is better to fly more before the tax comes. From the passenger point of view, they might over-estimate how big the effect on the price will be and figure that it is better to travel abroad a bit more while it is cheaper to do so. From the airline point of view, they might maximize the capacity from Sweden compared to for example Denmark while it is still cheaper to do so, or they might try to exaggerate the tax effect deliberately to be able to criticize it more effectively. When they then go back to previous levels they can then point to this reduction in passengers in order to get other forms of subsidies or compensations from the government. When this quarter is included in the pre-fit, the results are also higher (as expected) as can be seen in Table 1, but as the SCM does not allow for the estimation of confidence intervals it might very well be that the difference could be within this unobservable confidence interval. However, to address this potential issue more directly the start of the treatment time can also be adjusted and the potential anticipation effect can thereby be estimated. This is also recommended by Abadie et al. (2015) when an anticipation effect is suspected, i.e. one should then adjust the start of the treatment period to when the anticipation effects is assumed to have started. This is done in section 6, and it is clear that while the third quarter of 2017 is positive the effect is not significant. This should therefore not be a big issue in the main estimations, but for reasons of caution we still opt for the most conservative estimate, where the year before the tax introduction is exempt from the pre-fit, to be used as the main estimation of the tax effects.

While the results paint a clear picture so far, there are other concerns of potential biases that need to be addressed. One such bias is the potential for leakage. People living close to the border to a neighboring country, such as Denmark or Finland, could respond to the tax by traveling to the closest airport across the border instead of the closest airport in Sweden, as a result of the tax. While this definitely happened both before and after the tax was introduced, as some Swedes simply live closer to a Danish or Finish airport and it is the most convenient option for them, it becomes an issue if people who used to go to the Swedish airport now go to the airport across the border as a response to the tax. This is potential bias is especially important to address as both Denmark and Finland are apart of the synthetic control group in the passenger estimation, and as papers such as Borbely

(2019) finds evidence of a clear leakage effect to neighboring countries. However, when an estimation without Denmark and Finland (Norway is already excluded from earlier) is done, no sign of leakage can be found as the estimated effects are even higher than in the main estimation. This implies that there has not been any significant leakage to neighboring countries due to the tax.

6 Robustness checks

In this section, some robustness checks are considered to evaluate the plausibility of the results. The first check that is considered is a placebo treatment time test; the same method is used but the treatment time is changed to an earlier date. As no treatment happened at this time, no effect should be visible. As the data that is used in the original estimation is deseasonalized using variations up until 2017q2, two different versions of this robustness check is considered. The first one uses the same data as in the original estimation, and only changes the treatment time. This tests the validity of the SCM method on the data that is used in the main estimation. The second one makes a new deseasonalization using the seasonal variation up until the second quarter of 2014, which is the chosen placebo time, to deseasonalize the whole series. This instead evaluates whether the the method of deseasonalization itself introduces any bias in the estimated effects. Both of these estimations also give further insight into how trustworthy the tax estimates are; if the difference between the synthetic control group and Sweden after the tax is driven only by the tax we would expect no similar differences to happen any time between the placebo treatment time and the actual treatment time, in neither of the alternative estimations. The results of these two placebo treatment time tests are found in Figure 3.

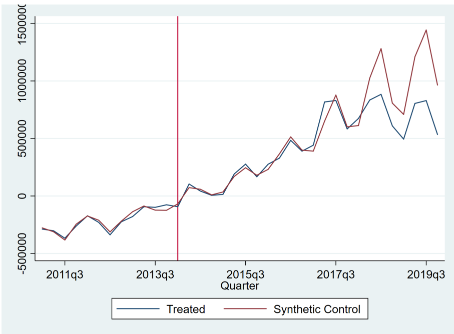

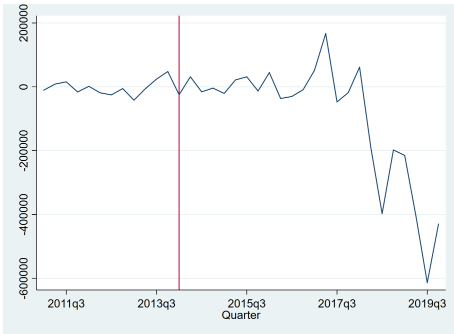

(a) Trend comparison, using the same deseasonalized data as in the main estimation

(b) Effects, using the same deseasonalized data as in the main estimation

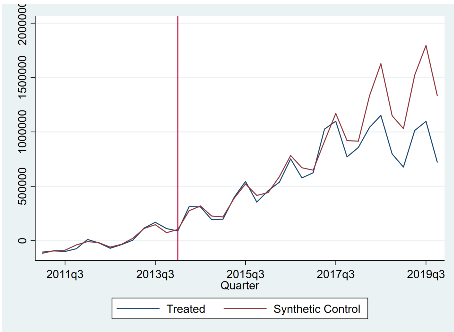

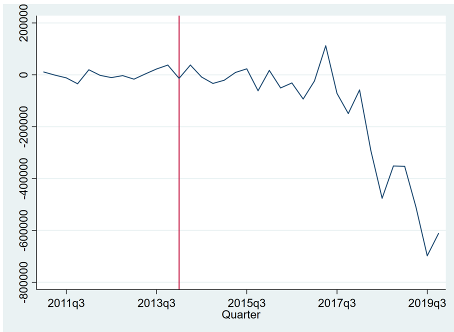

(c) Trend comparison, using data that is deseasonalized only until the new "treatment date"

(d) Effects, using data that is deseasonalized only until the new "treatment date"

Figure 3: Placebo time estimation with deseasonalized international passenger data calculated in two ways: one being the same as in the main estimation (i.e. deseasonalized until the actual treatment date) and one being deseasonalized until the placebo treatment date.

As can be seen, for both of these estimations the synthetic control group tracks Sweden between 2014 and 2018 almost as well as in the actual estimation. This therefore provide evidence for that the SCM does a good job in estimating a suitable control group and that the method of deseasonalization does not introduce any bias. The effect of the the tax, at the actual treatment time, is also clear in these estimations.

As previously mentioned, there could also be a bias present from a potential anticipation effect: that the announcement of the tax, which happened in the second quarter the year before the actual implementation of the tax, might have resulted in strategic behaviors by either the airlines or the passengers. The same SCM-procedure is therefore applied, but where the treatment time is set as the second quarter of 2017 instead of the second quarter in 2018. As can be seen in Table 5, the first significant effect is however found the first quarter the tax is actually implemented. The four quarters before the actual implementation is not significant. While the ocular examination of the development of international air passengers from Sweden indicates a potential anticipation effect, this deviation is however not significant. Still, as the there is a clear spike in Sweden the second quarter the year before the tax it is still good to be careful in choosing the main

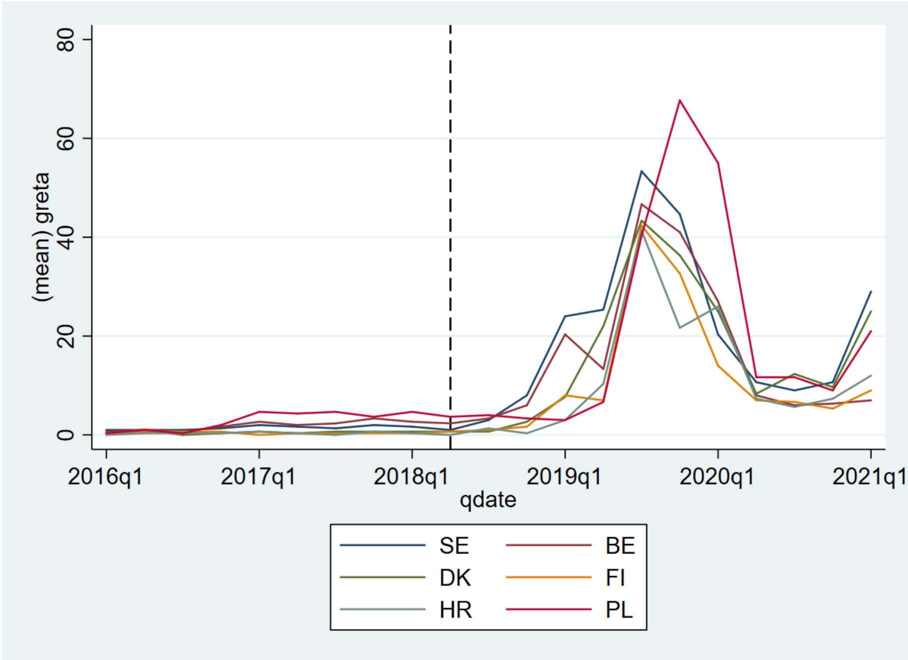

Figure 4: Google search intensity for "Greta Thunberg" in Sweden and the countries that are part of the synthetic control group in the main specification.

specification in the previous section. Choosing the most conservative estimate, as we do, is therefore wise.

Another issue worth addressing is the potential Greta effect: since Greta comes from Sweden and started to spread her message after the tax was introduced (especially the second tax year could be affected by this), it is possible that this effect has been bigger in Sweden than in other countries, thereby introducing a biased result. This very hard to test directly, but one way to gauge the potential differences in effects is to look at Greta's popularity in internet searches. In Figure 4 the Google trends statistics for the search "Greta Thunberg" is presented for Sweden as well as the countries in the (original) synthetic control group. As can be seen there is no evidence for a larger popularity in Sweden, and in particular the trends are changing in a similar fashion in all of the countries that are considered. As the Google trends data is an index relative to the country or region itself, the numbers are not directly comparable between the countries. However, as the change are so similar in all of the countries it still provides reassuring evidence for that the Greta-effect likely is not a big issue in this estimation. At the same time, similarity in trends in popularity does not ensure that the effects of the popularity is the same in all countries.

Another way to assess whether Greta has had an increased effect is to estimate the effect on passenger number using the similar taxes that was introduced in Norway and Austria, especially as the tax effect in Sweden seems to grow for every quarter after the tax introduction. If the effects follow a similar pattern, it is less likely that Greta has

Table 5: Formal examination of a potential anticipation effect from the announcement of the introduction of the tax

	(1) Passengers (std. p-value)
2017q2	147140.9
	.1428571
2017q3	34244.59
	.6666667
2017q4	-83874.55
	.2380952
2018q1	-54818.02
	.5238095
2018q2	-236889.6
	0
2018q3	-350491.9
	0
2018q4	-340914.6
	0
2019q1	-398167.1
	0
2019q2	-507309.8
	.047619
2019q3	-604190.2
	.047619
2019q4	-600436.3
	0

introduced any substantial bias in the case of Sweden. These estimations are done in a similar fashion as in the main specification: the data is deseasonalized and the SCM is applied using the same donor pool of countries for potential controls as in the main estimation. Sweden is however also included in the donor pool, which is particularly important to achieve an acceptable pre-fit for Norway, but the time period is restricted to the time before the tax was introduced in Sweden, to avoid bias at the end of the estimation periods. As can be seen in Figure 5 , the effects of the tax in these countries look very similar to the effects in Sweden: the effects seem to be small initially and then increase over time, just as for Sweden. This provides further evidence against any specific Greta-effect being responsible for the increased effects over time in Sweden. At the same time it provides further evidence for that the introduction of the taxes, in all these countries, coincide with an increased environmental awareness, and that the introduction of an environmentally motivated tax on air travel provides a strong signal from the government in the respective country that has effects beyond the price effect. This is especially so as the this type of pattern can not be solely explained by a price effect, as the tax is constant in all cases.

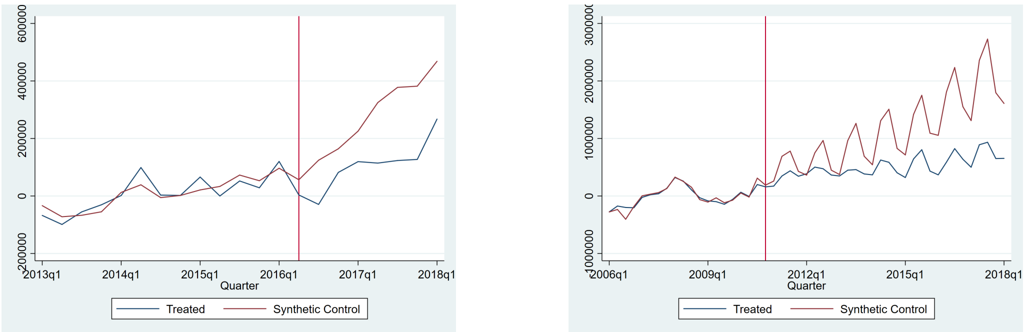

(a) **(b)**

Figure 5: Effect of Norwegian (a) and Austrian (b) tax on de-seasonalized international air travel

7 Conclusion

In this paper, the effects of the Swedish aviation tax on the passengers numbers and price index of air travel are estimated using the synthetic control group method. In addition, the price elasticity of Swedish international air travel is estimated using web-scraped price data and an instrumental variable approach to account for endogeneity issues. The tax effect on the number of departing international passengers is big and robust, starting at a decrease of around 4 percent which then increases to a decrease of around 10 percent. The tax effect on the price index of international air travel is not as robust, but indicates an increase in prices the first quarters after the introduction of the tax, the effect being as much as 10 percent the second quarter after the introduction, which then decreases and being less significant the second year after the tax. Given the estimated price elasticity of -0.76, the estimated price effects can explain the reduction in passengers the first three quarters. Since the effect on passengers increases as the effect on price decreases for the last 4 quarters of the period that is being studied, this indicates that there are additional behavioral effects that are present. One likely candidate is that the introduction of the tax increased the environmental awareness of the passengers even further through a symbol effect of the tax. While no direct evidence can be found for a potential “Greta-effect”, whose rise to fame coincides with the last 4 quarters of the period of study, it is still also possible that this environmental awareness is fueled even more buy her messages. In general, the pro-environment messages are strong in Sweden, and have been gaining more and more traction over the past years, and it is indeed possible that this explains the increased effects in the second year after the introduction of the tax, especially as the price effects are reduced after the initial quarters after the tax. Given the timing of the relative reduction in international air travel, it is however clear that the tax had a large effect on the demand for air travel, especially as the passenger effects are indeed very robust and since the potential biases have been addressed. The main take away is therefore that in the case of Sweden, an aviation tax was very successful in reducing the number of international passengers. The results from this paper can therefore be used by policy makers both in Sweden, but also in countries that are similar to Sweden.

While this paper makes an important contribution to the relatively sparse literature on the effects of aviation taxes, there are still several areas of interest for future research. Given that there seems to be a behavioral effect in addition to the price effect of the tax, one interesting research topic is to compare the effects of a newly introduced tax with the effects of an increase in already existing taxes. Intuitively the effect should be larger in the initial introduction of a large part of the tax’s effect is due to a symbol effect of the tax, especially since a newly introduced tax likely gets more media attention than an increase of an already existing tax. In general, stronger evidence on the symbol versus price effects of an environmental aviation tax is needed in the literature. In addition, given that the price effect estimations in this paper are far less robust than the passenger effect estimation, further research on the effect on prices are also needed. Ideally, one would have actual price data for at least a year before the introduction of the tax, something that we did not have in this paper as we only had three months of pre-tax data. Finally, since the effect of aviation taxes seems to differ between countries it is important to also evaluate future aviation taxes to get a broader picture of its effects.

Appendix

A.1 Tax zones

Table 6: Countries in the first two tax-zones
(the rest of the countries are in the third tax-zone)

Tax zone 1	Tax zone 2
Albanien	Afghanistan
Andorra	Algeriet
Belgien	Armenien
Bosnien och Hercegovina	Azerbajdzjan
Bulgarien	Bahrain
Cypern	Burkina Faso
Danmark	Djibouti
Estland	Egypten
Finland	Elfenbenskusten
Frankrike	Eritrea
Grekland	Etiopien
Irland	Förenade Arabemiraten
Island	Gambia
Italien	Georgien
Kosovo	Guinea
Kroatien	Guinea-Bissau
Lettland	Irak
Liechtenstein	Iran
Litauen	Israel
Luxemburg	Jemen
Makedonien	Jordanien
Malta	Kanada
Moldavien	Kap Verde
Monaco	Kazakstan
Montenegro	Kirgizistan
Nederländerna	Kuwait
Norge	Libanon
Polen	Libyen
Portugal	Mali
Rumänien	Mauretanien
San Marino	Marocko
Schweiz	Niger
Serbien	Oman
Slovakien	Pakistan
Slovenien	Palestina
Spanien	Qatar
Storbritannien	Ryssland
Sverige	Saudiarabien
Tjeckien	Senegal
Turkiet	Sudan
Tyskland	Syrien
Ukraina	Tadzjikistan
Ungern	Tchad
Vatikanstaten	Tunisien
Vitryssland	Turkmenistan
Österrike	USA
	Uzbekistan

A.2 Donor weights in the main tax effect estimations

Table 7: Donor weights

	Passengers	Price
AT	0	.068
BE	.035	0
BG	0	0
CY	0	0
CZ	0	0
DK	.325	0
EE	0	.13
FI	.255	.12
HR	.188	0
HU	0	0
IE	0	0
IS	0	.292
LT	0	.132
LU	0	.03
LV	0	0
MT	0	0
PL	.197	.048
PT	0	0
RO	0	0
SI	0	.142
SK	0	.038

A.3 Development of predictor variables

In figure 6, the development of the predictor variables in Sweden and countries that are part of the synthetic control group for the estimation of international passengers are presented.

(a) Development of GDP/capita in Sweden and synthetic control group countries

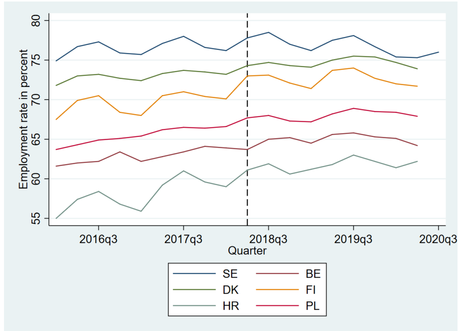

(b) Development of employment rate in Sweden and synthetic control group countries

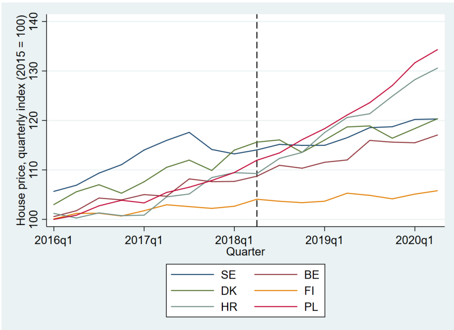

(c) Development of house prices in Sweden and synthetic control group countries

(d) Development of price of air travel in Sweden and synthetic control group countries

Figure 6: Examination of developement of predictor variabels in the syntehtic control group countries in comparison to Sweden.

References

Abadie, A., A. Diamond, and J. Hainmueller (2010, 6). 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, 2). Comparative Politics and the Synthetic Control Method. *American Journal of Political Science 59*(2), 495–510.

Berry, S. and P. Jia (2010, 8). Tracing the Woes: An Empirical Analysis of the Airline Industry. *American Economic Journal: Microeconomics 2*(3), 1–43.

Borbely, D. (2019, 8). A case study on Germany's aviation tax using the synthetic control approach. *Transportation Research Part A: Policy and Practice 126*, 377–395.

Brons, M., E. Pels, P. Nijkamp, and P. Rietveld (2002, 5). Price elasticities of demand for passenger air travel: a meta-analysis. *Journal of Air Transport Management 8*(3), 165–175.

Falk, M. and E. Hagsten (2019, 1). Short-run impact of the flight departure tax on air travel. *International Journal of Tourism Research 21*(1), 37–44.

Forsyth, P., L. Dwyer, R. Spurr, and T. Pham (2014, 2). The impacts of Australia's departure tax: Tourism versus the economy? *Tourism Management 40*, 126–136.

Gayle, P. G. (2004). Does Price Matter? Price and Non-Price Competition in the Airline Industry. *Econometric Society 2004 North American Summer Meetings, Econometric Society No. 163*.

Granados, N., A. Gupta, and R. J. Kauffman (2012). Online and offline demand and price elasticities: Evidence from the air travel industry. *Information Systems Research 23*(1), 164–181.

Hsiao, C. Y. and M. Hansen (2011, 11). A passenger demand model for air transportation in a hub-and-spoke network. *Transportation Research Part E: Logistics and Transportation Review 47*(6), 1112–1125.

Kopsch, F. (2012, 5). A demand model for domestic air travel in Sweden. *Journal of Air Transport Management 20*, 46–48.

Mayor, K. and R. S. Tol (2007, 11). The impact of the UK aviation tax on carbon dioxide emissions and visitor numbers. *Transport Policy 14*(6), 507–513.

Mumbower, S., L. A. Garrow, and M. J. Higgins (2014, 8). Estimating flight-level price elasticities using online airline data: A first step toward integrating pricing, demand, and revenue optimization. *Transportation Research Part A: Policy and Practice 66*(1), 196–212.

Peng, B., H. Song, G. I. Crouch, and S. F. Witt (2015, 9). A Meta-Analysis of International Tourism Demand Elasticities. *Journal of Travel Research 54*(5), 611–633.

Seetaram, N., H. Song, and S. J. Page (2014, 7). Air Passenger Duty and Outbound Tourism Demand from the United Kingdom. *Journal of Travel Research 53*(4), 476–487.