Tax Flights

Koleman Strumpf*
University of Kansas School of Business

Preliminary results only

November 30, 2012

Abstract

Tax evasion is difficult to measure, since evaders try to avoid detection and counter-factual behavior is hard to establish. This paper considers evasion in an environment where these two issues can be overcome. Aircraft are taxed as personal property in some American states. Taxes are owed if the plane is hangared in the state on one specific date. Strategic plane owners may try to evade the tax by flying to a non-taxing jurisdiction just before this date and returning shortly thereafter. I assess such "tax flights" using a database of about twenty million trips covering general aviation flights in the United States during the period 2004 to 2009. For each flight I know the time, location of the arrival and departure airport, the address of the owner, and the type of plane. I match this to a database of local tax rates and valuation of planes to measure the potential tax bills. To establish the counter-factual flying behavior, I exploit variation in tax policy (at both the state and local level), exemptions for certain classes of planes, type of plane, tax valuation method, and tax date. Preliminary results indicate the presence of tax flights. In later revisions I will estimate this elasticity by demographics, such as income, which can potentially inform both theory and practice. To validate the results, I will see whether planes on tax flights are actually missing from tax rolls and also employ various placebo tests such as considering the impact of taxes on exempt plane flight patterns and quasi-experimental tax rate changes applying only to certain classes of planes.

*I would like to thank Dan Fetter, Heidi Williams, Jeff Zabel, and participants at DePaul University and the Lincoln Institute Urban Economics and Public Finance Conference for comments. The FAA shared archive copies of the Aircraft Registry, Maponics LLC furnished geocoded ZIP+4 databases, and FlightView Inc provided general aviation flight logs. A portion of this paper was written while I was a Visiting Scholar at the Management & Strategy Department at the Kellogg School of Management. Financial support from the Lincoln Institute is gratefully acknowledged.
Correspondence should be sent to 1300 Sunnyside Avenue Lawrence, KS 66045 or cigar@ku.edu.

1 Introduction

A central issue in public economics is the extent to which individuals or firms evade taxes. It is typically difficult to quantify such evasion, since it is hard to observe (evaders hide their actions) or establish the counter-factual (what behavior would have been like in the absence of taxes). This paper considers an application, the property taxation of general aviation (GA) aircraft, in which these issues might be overcome. These taxes are levied in some states and are based on the plane's location on a specific date referred to as the assessment date.

Strategic plane owners might try to evade the property tax by flying their plane to a non-taxing jurisdiction just before the assessment date and return shortly thereafter.[1] Such *tax flights* could plausibly succeed since planes are mobile and tax authorities rarely have a complete database of all planes in their jurisdiction (in contrast to other property such as homes or autos).

Precisely measuring tax evasion is possible in this environment. The researcher can observe the underlying behavior because the Federal Aviation Administration (FAA) tracks and caches GA flights. The counter-factual of how many flights there would be around the assessment date in the absence of taxes can be established using variation in tax policy (at both the state and local level), in exemptions for certain classes of planes, in type of plane, in tax valuation method, and in the assessment date. Netting out the counter-factual behavior from actual flights around the assessment date gives a measure of tax flights.

In this paper I use a database of about twenty million trips covering GA flights in the United States during the period 2004 to 2009. For each flight I know the time, location of the arrival and departure airport, the address of the owner, and the type of plane. I match this to a database of local tax rates and valuation of planes to measure the potential tax bills. Preliminary results indicate the presence of tax flights. As a verification check on a data subset, I find that the tax flight planes are disproportionately missing from local tax rolls.

I build on the large literature which empirically measures tax evasion (see the summary in Andreoni, et al 1998; Slemrod and Yitzhaki, 2002; Slemrod, 2007).[2] Re-

[1]Senator Claire McCaskill appears to have used such a strategy to evade \$300,000 in property taxes over four years on a plane she co-owns (Scott Wong and and John Bresnahan, 21 March 2011, "McCaskill to pay back taxes on plane," *Politico*).

[2]Tax evasion is formally defined as willful actions which result in the illegal underpayment of

cent papers have relied on range of approaches to calculate the extent of tax evasion including changing regulator detection strategies (Marion and Muehlegger, 2008), exploiting behavioral responses (Chetty et al, 2009), examining tax code discontinuities (Kleven and Waseem, 2011; Chetty et al, 2011), developing novel data sets (Merriman, 2010) and inferring third party information (Artvanis, et al 2011). A challenge for these papers is to verify the accuracy of their measure, which is difficult because the underlying behavior is unobserved. I can more directly establish evasion occurs, since I explicitly observe the hidden behavior of tax non-compliers. In particular I verify whether individual agents my estimates imply are evading taxes are actually missing from tax rolls. Because of unique aspects of the environment I study and the very rich data set I have, I also am able to consider additional issues beyond the focus of previous work. In particular I build on two important recent papers, Fisman and Wei (2004) and Kleven, et al (2011). Fisman and Wei (2004) measure tariff evasion in the Hong Kong-to-China trade by comparing export and import statistics. They find exports exceed imports for goods facing high tariffs, with some of the missing exports reappearing as imports of lower tariff goods. I use a similar identification strategy in my approach, and in addition I can observe actual evasion at the decision-maker level and can use completely non-taxed planes to help establish the counter-factual. Kleven, et al (2011) measure evasion at the taxpayer-level using a random audit in Denmark. I also analyze individual taxpayers, with the addition that the potential savings from evasion is typically larger and there are business as well as individual decision-makers.

The paper also adds to the literature on the aviation industry. Most papers here focus on commercial carriers, and address issues such as the impact of hubbing on firm performance (Mayer and Sinai, 2003), the impact of deregulation (Winston and Morrison, 1995), evidence of price discrimination in ticket prices (Borenstein and Rose, 1994), response to potential entry (Goolsbee and Syverson, 2008), rules for optimal airport congestion pricing (Brueckner, 2002), or factors influencing vertical integration (Forbes and Lederman, 2009). This paper has a different focus, looking at issues related to public economics rather than industrial organization. Also I study a different segment of the industry, general aviation, which allows me to investigate differences between private and commercial owners which cannot be evaluated using

taxes. In contrast tax avoidance involves legal tax mitigation strategies. The behavior in this paper is legally murky, but I will refer to it as tax evasion.

scheduled airline data.

2 Background

2.1 Institutional Framework

This paper focuses on GA aircraft which includes almost all civil aviation besides airlines. It includes both commercial and non-commercial aircraft, and it includes a wide range of plane types including reciprocating (piston) engines, turboprops, light jets, and experimentals. GA can have individual or firm owners, and they span from inexpensive kit models to multi-million dollar jets. There are over 13k GA airports in the US, and 350k GA aircraft registered with the FAA (about a third of these planes are inactive).

Figure 1 maps state tax policies on GA aircraft (The Data Appendix contains a list of sources used to generate the stylized facts in this section). Eighteen states allow local governments to levy some form of personal property tax on these planes. While most taxing states are in the south or west, there are non-taxing states in all regions (in 2010 forty percent of GA traffic involved taxing states). Among taxing states, twelve tax all aircraft, five tax just business-owned aircraft, and one taxes just personal-owned aircraft. The taxing states assess planes on a single date, which is 1 January in sixteen cases and other dates in two others. In seventeen of the states there is a uniform method of determining assessed values (a fraction of current retail or wholesale price, a depreciation schedule based on purchase price, and other permutations) and one state allows each county to pick their own method. Several states also have a variety of exemptions for particular planes (such as planes older than a certain age or planes used in agriculture). States also have a variety of tax situs, with sixteen taxing planes where they are located and the others based on where the owner is located.

The property tax system is locally administered (Unlike with autos, there is no state registry of all planes. The FAA keeps a registry which it updates semi-monthly). While the state sets the basic rules as described in the last paragraph, counties are in charge of collecting the tax. Most tax officials appear to devote little time or expertise to aircrafts.[3] A reason for this is few counties have specialists in aircraft

[3]A graphic example of this may be found in Ryan Kath (2011), "Investigation

taxes, and the division which typically administers it is primarily focused on real property such as homes. Still, some counties have requested a list of planes hangared at local airports on the assessment date (California and Nebraska statutes require airports or hangars to report the list of based planes on the assessment date). This appears to be the main form of detection, so a tax flight away from the airport just before the assessment date would be a simple means of evasion. That is, the plane is unlikely to be detected though the flight does not remove the legal obligation to pay taxes. The tax flights might be unsuccessful when local tax authorities engage in more sophisticated strategies, such as consulting online sources listing recent flight activity by plane.[4]

The mechanics of aircraft property taxes typically parallel those on other property. The tax owed on a particular plane is the product of its assessed value and the overall set of rates. The assessed value is based on the state system of valuation applied to the specific assessment date. The rate is the sum of all those from taxing jurisdictions, which may include the state, county, municipality, school district, and special districts. These rates are typically adjusted each year. A key difference from other forms of property taxation is that no bill is typically sent out, but rather owners are responsible for submitting forms along with payments.

An important question is what happens to a plane owner who is found to have evaded taxes. There do not appear to be clear rules on this but from extensive discussions with local and state tax authorities (see Data Appendix) it appears that the owner typically must pay all back taxes plus a multiplicative factor which is proportionate to the unpaid taxes. That is the payment is proportionate to the amount of taxes which have been evaded. This condition will be used in the next sub-section.

2.2 Simple Model of Tax Evasion

Consider an owner who is deciding whether to evade property tax payments on his plane. This is a version of the standard Allingham-Sandmo-Yitzhaki type model in

finds dozens of plane owners not paying taxes, costing local governments big bucks." http://www.nbcactionnews.com/dpp/news/local_news/investigations/investigation-finds-dozens-of-plane-owners-not-paying-taxes,-costing-local-governments-big-bucks-may2011swp.

[4]Tax authorities can also consult plane registries. But these list where the owner, but not the plane, are located. This information is not as useful for enforcement in the majority of taxing states which use plane location as the basis for tax situs.

which the choice variable is discrete and where the only a portion of income is taxable. Suppose the plane has assessed value B and faces a property tax rate of t. If he does not evade he pays taxes of tB. If he evades, he is caught with probability p and must pay a penalty $\Delta > 1$ on the understated taxes, and if he is not caught then he pays no taxes. A risk averse owner with other income I will evade if,

$$Evade \leftrightarrow (1-p)U(I) + pU(I - \Delta tB) > U(I - tB) \tag{1}$$

where $U(\cdot)$ is the utility function with $U' > 0$ and $U'' > 0$. The left hand side of the inequality is the expected utility of evading, with the first term representing the case where the owner is not detected and the second term is the case where he is detected.

Under this framework, the following comparative statics hold. The propensity to evade is decreasing in the probability of detection (p) as well as in the penalty (Δ). The other terms have an ambiguous effect, e.g. both the benefit (avoided tax) and cost (penalty) of evasion are increasing in the plane value. In practice p is quite small in which case the propensity to evade is increasing in the tax rate (t), the value of the plane (B) and is decreasing in income (I). In principle all of these implications are testable. However I do not have data about the first two points, so in the empirical application I will focus on the relation between evasion and tax rate, plane value and owner income.

2.3 Identification

The key question is how much flight activity, presumably wasteful, does this tax system induce. The extent of tax evasion can be measured from several sources of variation:

(i) taxing versus non-taxing jurisdictions: one can compare flights in states which allow local governments to levy property taxes with those in non-taxing states;

(ii) tax rates and assessment methods: in states which allow taxes, local governments vary in both the rates they apply and their methods of setting assessed values;

(iii) plane types: some planes are more valuable than others, and as such face different potential tax burdens if they do not evade;

(iv) special exemptions: some states only allow taxation of certain kinds of planes, such as business-owned, non-business owned, or those less than a certain age;

(v) a natural experiment (West Virginia effectively made business planes exempt in 2009 while previously all planes were taxed).

For reference Figure 6 overlays tax units in the taxing states (in red) on a map of all airports in the U.S.

The goal is to see the change in behavior of the treated group (plane owners facing a property tax) relative to a control period (non-taxing period) and relative to control planes (plane owners not facing the property tax). The main specification to be estimated is,

$$\begin{aligned} Flights_{igt} &= \beta_1 TaxTime_{gt} \times TaxState_g \times TaxBill_{igt} + \beta_2 TaxTime_{gt} \times TaxState_g \quad (2) \\ &+ \beta_3 TaxState_g \times TaxBill_{igt} + \beta_4 TaxTime_{gt} \times TaxBill_{igt} \quad (3) \\ &+ \beta_5 TaxTime_{gt} + \beta_6 TaxState_g + \beta_7 TaxBill_{igt} \\ &+ \alpha + X_{igt}\gamma + \epsilon_{igt} \end{aligned}$$

where i = plane, g = geographic location (state or local government), t = date, $Flights$ = a measure of flight activity, $TaxTime$ = an indicator for assessment time in that state, $TaxState$ = an indicator for a state that taxes planes, $TaxBill$ = plane i value in g times the tax rate in g at time t, X = controls such as weather (some specifications will also have μ = fixed effects at the plane, geography or time-level, though this will eliminate some of the Tax terms). The key parameters are β_1 and β_2, which measures how flight activity changes in a taxing state around the assessment time and whether this effect changes with the tax rate. The estimation section contains results from simpler specifications which omit many of the terms to facilitate in interpretation.

The specification can be thought of as either a regression discontinuity or difference-in-difference design. From either perspective, we can think of comparing planes located near the border of a taxing and non-taxing state, comparing planes which are subject to the tax with those that are exempt, or comparing a taxed plane's flights just before/after the assessment date to further off periods. In the case of the West Virginia law change, we can compare business plane flights in the state after the exemption was introduced to previous years, compare business plane flights to personal plane flights before and after the exemption, and compare these to comparable differences in other states. The key in all these cases is that there is distinct treatment group (non-exempt planes in a taxing state during the tax period) and control group (otherwise). In addition, there are continuous treatment variables, such as the tax rate (which changes over both jurisdictions and over time within a jurisdiction) or taxable value of the plane (which varies across plane type, over time within a jurisdiction, and between jurisdictions due to differences in assessment systems).

A final issue is concerns about endogeneity. It may be that unobserved factors of flight activity (ϵ in the specification) are correlated with the tax bill. For example, plane value might influence flights. But fixed effects largely account for this possibility. Another possibility is that governments take into account tax flights when they set tax rates (we only have to be concerned about tax rate *changes* due to the fixed effects). But we have already seen that governments do not closely monitor airplanes so this is unlikely. In addition, this would be hard to implement since as discussed in the next section the same property tax rate is used for other forms of personal and sometimes real property and the overlapping taxing jurisdictions would have to coordinate their rates. Still it is possible to directly account for endogenous tax rates. First, I can refine the $TaxBill$ term in the specification to just be an indicator of whether this is a taxing state so the variation in rates is eliminated. Second, I can instrument for the tax rates using characteristics of the property tax system (the timing of reassessment or exemptions up to certain property values) which are primarily set based on real property.

3 Data

3.1 Sources

There are several data sets which have to be integrated for the analysis. The first step is to assemble a database of annual aircraft tax rates. Planes are taxed as tangible personal property, and the rate is typically the general personal rate. An overlapping set of jurisdictions may levy such taxes, including the state, county, municipality (city, borough, township and other sub-county political sub-divisions), and school district (unified, secondary and elementary).[5] While all counties may tax planes, each state has different rules on which of the other government types are permitted to tax.[6] Figure 2 displays the tax units for Texas as an illustration. The tax rate database draws from the Lincoln Institute's *Significant Features of Property Tax* (2010), which lists rates at the county and sometimes sub-county level. A variety of state-specific sources discussed in the Data Appendix is then used to fill in the remaining rates. Figure 3 shows an example of the rates for Texas in 2009.

The second step is to determine the taxable values of each plane. This is based on *Aircraft Bluebook Historical Value Reference* (2010) and is discussed in the next sub-section.

The third step is to associate with each plane the set of taxing jurisdictions, and thus the tax rate and assessed value. Initially various addresses have to be geolocated (determine their longitude and latitude). As described in the last section, some states tax planes based on their location and others base it on the owner's location. Plane locations are based on the airport coordinates in the FAA's *Form 5010: Airport Master Record* (2010) and additional sources listed in the Data Appendix. The owners' addresses are listed in the FAA's *Aircraft Registry* (various years).[7] Through special arrangement with the FAA, I have copies of this file for each month over the period March 2004 to July 2009. Each file is geocoded using a three step process discussed in the next sub-section, and then the coordinates are matched to taxing jurisdictions using the ArcGIS software package and the Census' *TIGER/Line Shapefile* (various years). Every location in the United States is located in exactly one state, county, county subdivision, and school district (unified or elementary/secondary). A subset are also located in places. All municipalities are listed as either a county subdivision or place. Two examples of the output are mapped in Figures 4 and 5 (Figure 6 overlays tax units in the taxing states on the last map). The next sub-section discusses this address geocoding procedure in more detail and also its current status.

[5] In many states single purpose special districts can also levy taxes, but it is not possible to geographically locate all such district and to match them to addresses as described later in this section. I therefore calculate the average rate for each category of special district (safety, fire, sanitation, water, etc) in each county, and then add the sum of these averages to the county rate.

[6] Among states allowing plane taxes, only Virginia prohibits school districts from levying a property tax.

[7] All plane owners must register their planes with FAA once every three years. These registrations are the basis for the *Aircraft Registry*. Note that the database includes many inactive planes, since the FAA does not expunge all planes which have not re-registered.

The final step is to generate a database of plane flights. A log of GA flights in the US for the period January 2004 through July 2009 come from FlightView Inc. These data are generated in the course of normal flight activity when a pilot registers his flight plan with the FAA. The FAA sends a live feed of the flight information to authorized vendors under the Aircraft Situation Display to Industry (ASDI) program. Vendors, such as FlightView, translate the feed into a usable format and remove anomalies (FAA, 2009 provides background on the ASDI program).[8] The final data include the date, the tail number, the aircraft type, the arrival/departure time and airport, and distance between these airports. There are 210k unique planes and 24m flights.

3.2 Complications

Some of the data is not yet available in a form amenable to empirical analysis. The three issues listed below will be completed in the next revision of the paper.

Table 1 highlights some of the issues with the local tax rate data. There are several thousand tax units in Texas, Nebraska, Kansas, while there are unusual circumstances in Nebraska, Virginia and Louisiana. California does not have a centralized database of tax rates (according to its Board of Equalization), so county averages will be used.[9]

A second issue is to determine the taxable value for each airplane. Each plane will be matched to the *Aircraft Bluebook Historical Value Reference* (2010) using its manufacturer, model and manufactured year. The values listed in the *Bluebook* are updated quarterly. The main difficulties here are accounting for the various assess-

[8]There are two sets of flights which are omitted from this feed. First, a plane owner can select to block his plane from either the general FAA feed or from a specific ASDI vendor database (the procedure is discussed in NBAA, 2010). Second, flight logs are only required under instrument flight rules (IFR) while a pilot can instead fly under visual flight rules (VFR) when weather conditions are favorable and the plane does not fly into certain restricted airspaces. A concern is that pilots may strategically utilize one of these options as a method to evade property taxes. There are reasons to doubt these possibilities. First, the blocked list is rather small and is largely composed of planes whose owners are public figures or large corporations (Michael Grabell and Sebastian Jones, 8 April 2010, "Off the Radar: Private Planes Hidden From Public View," *ProPublica*; Mark Maremont and Tom McGinty, 21 May 2011, "For the Highest Fliers, New Scrutiny," *Wall Street Journal*). Second, the proportion of VFR flights actually decreases in the period just before and after an assessment date in taxing states.

[9]Proposition 13 limits property tax rates in California to one percent, except when a super-majority of voters approve additional levies for school bonds or facilities. In practice this means there is therefore only small differences in tax rates in the state, and so using county averages omits relatively little variation.

ment systems used in different states (based on retail value, on wholesale value, on depreciation schedules, or some other system), the prevalence of kit and experimental planes, and adjusting for modifications in each aircraft.

The final issue is geocoding addresses. Figure 7 is the flow chart of the process. In the first step the FAA's *Aircraft Registry* address files, which contain over three hundred thousand records, must be converted from pdf to text format. The next step in the second row shows how coordinates for each address are obtained. The full street addresses are matched to a year-specific database in ESRI ArcGIS (various years), then the zip codes from unmatched addresses are compared to nine-digit zip databases from Maponics (2010) and the USPS (2010). The last step, shown in the remaining rows, is to locate the relevant taxing units (state, county, place, county subdivision, unified/secondary/elementary school district and additional units in certain states) for each of the matched addresses. Each type of unit must be matched separately using an ArcGIS spatial join to the Census' *TIGER/Line Shapefile* (various years). Roughly 85% of the addresses can be geolocated in this fashion. This process takes roughly a week of processing time for each set of data, and this in not yet complete since there are about sixty sets of addresses (corresponding to each monthly FAA registries). This entire process must be repeated for the list of airport locations.

4 Preliminary Results

4.1 Motivating Graphs

Before turning to the estimates, it is helpful to visualize the data. Figure 8 shows the weekly number of GA flights for each year between 2004-2009 (only the first half of 2009 is available). A clear seasonal pattern is apparent with a peak during the summer months and a rough in the winter months. This is important for the estimates since the assessment date for sixteen of the eighteen taxing states is 1 January, which is near the trough. There is also a sharp drop in traffic around week 27 which includes the 4 July holiday and is near the assessment date for another state. There is also a drop in traffic in the last three years, likely due to the deep recession at that period. These temporal patterns point out the importance of including both week and year fixed controls in the estimates.

The remaining figures provide some preliminary evidence of tax flights by considering state-level flight patterns. If tax flights occur, then in taxing states there should be a dip in the number of planes located at their "home" airport just before the assessment date and this number should revert right after the assessment date. In non-taxing states, there should be no such dip except after accounting for seasonal flight patterns. A second implication of tax flights is that taxing states should have an increase in out-of-state traffic just before the assessment date, and an increase in into-the-state traffic just after the date. Non-taxing states should have the opposite pattern as planes evading taxes fly in and then leave. Comparing the trends away from the assessment period allows us to see whether the non-tax states serve as a suitable control group.

Figure 9 examines home airport patterns. Since it is not completely clear how to determine where a plane is based, I consider three separate definitions of the home airport: (i) the one where it spends the most time on the ground between flights; (ii) the one where it has the most arrivals plus departures; (iii) the one where it has the most round-trips (flights in which the arrival and destination airport are identical).[10] For each of these definitions, I calculate the proportion of active planes which are at their home airport at least once in each week of the year. I then divide the planes by whether their home airport is in a taxing state or not. Figures 9 show the results for the weeks just before and after the typical 1 January assessment date.[11] There is a comparable pattern in all cases. The taxing states see a sharp drop in home airport presence right before the assessment date and then a near reversion to their previous level in the weeks after the assessment. While this is consistent with tax flights, another explanation is that owners are going on an end of the year vacation. The non-tax states provide a control for this. While there is a dip and reversion in home airport presence in non-tax states in this period, it is far smaller and smoother than with the tax-states. Note that aside from the weeks further before or after the assessment date, the two series trend together suggesting that the non-taxing states are a suitable control group.

Figures 10-11 show that inter-state flight patterns are also consistent with tax

[10]The home airport may be undefined (some planes have no round-trips) or ambiguous (multiple airports can have the same number of summed arrivals and departures). The results discussed below are robust to different approaches to dealing with these cases (e.g. omit planes no unique home airport or include just one of the airports).

[11]The figures in this section assumes the assessment date is in week 1, which it is for sixteen of the eighteen taxing states. The other two states are omitted from the taxing group for the purposes of this figure.

flights. For both taxing and non-taxing states, the number of out-of-state and in-state flights closely track each other in most weeks, but they deviate in the weeks around the 1 January assessment date. In taxing states in the week before the assessment outbound flights exceed inbounds and the reverse holds just after the assessment period. For non-taxing states the opposite pattern holds, with a higher level of inbound flights before the assessment date and more outbounds afterwards (note that the asymmetry need not hold since the planes flying out of or into tax states could be coming from other tax states). Figure 12 show the same wedge in the neighborhood of the assessment date is evident in each year between 2005 and 2008. Tax flights are consistent with these figures, since it implies planes fly out of taxing states just before the assessment date and return shortly thereafter.

While the graphs here are suggestive of tax flights, because individual planes are not followed it is not conclusive. In particular I have to show that it is the same planes which are making the outbound and then inbound flight (it is possible that the outbound planes stay out of the state and a separate set of planes fly in to replace them). Moreover, the graphs might understate the extent of tax flights since state-level aggregation eliminates much of the variation in the data: tax rates, plane valuation, and exemption status. The estimates in the next sub-section address each of these points.

4.2 Preliminary Estimates

Table 2 shows how the sample of flights is constructed. Starting with the full list if 24.5m flights, about 0.5m are eliminated due to issues with matching to airports. The resulting set of 24m flights will be referred to as the most aggressive sample. Another 3m flights are removed for planes in which aircraft information is unavailable, and the sample of 21m remaining flights is the aggressive sample. Finally, another 0.5m flights are eliminated if there are consistency issues with the flight history, such as an departure time preceding the arrival time of the plane's most recent flight. This sample of 20.5m flights will be referred to as the conservative sample. Table 3 presents summary statistics which detail flight numbers for various subsets, counts of plane type characteristics, and a summary of tax rates.

Following the simplified model in Section 2.2, the tax flights hypothesis has predictions about the propensity of a plane owner to fly his plane at a particular time.

Analysis at the flight level is inappropriate since non-flying planes are ignored and highly active planes are over-represented. Instead the raw data are transformed to the week-plane level. Week-planes are included if the plane is actively flying or if there is flying activity at both an earlier and later date. This yields an unbalanced panel of roughly 26m flight-weeks (210k planes $\times$291 weeks minus weeks before/after the plane enters the sample). The analysis will focus on various weekly binary flight measures from the home airport using the hours on the ground measure discussed in the last sub-section (the results are comparable using either the flights or roundtrip measure).

Table 4 presents logit estimates of the propensity to be at the home airport at the end of the week (I am unable to estimate the full version of (2) since the data underlying $TaxBill$ are not available. In particular: (i) local tax rates by year are missing so instead I use county median property taxes per assessed values over 2005-2009 which are described in the Data Appendix; (ii) assessed values are missing since blue book values are not yet online, so instead I will use aircraft engine type as a proxy.). Column (1) shows that planes whose home airport is in a taxing state more likely to be at their home airport, though this effect is rather small (the odds of being at the home airport, relative to a non-taxing state airport, increase by a factor of 1.11). Column (2) adds terms involving an indicator $PreTaxTime$ which in a taxing state takes on a value of one in the week before the assessment date and zero otherwise and is similarly defined for a non-taxing state using the typical 1 January assessment date. The negative term on the interaction $Tax\ State \times PreTaxTime$ shows that planes in a taxing state are absent from their home airport just before the assessment date relative to other periods and to non-taxing states (the odds ratio is 0.74). This is consistent with tax flights, as owners fly away from their home airport just before their planes are assessed. Note that consistent with the graphs in the last sub-section there is an important seasonal effect, the negative parameter on $PreTaxTime$, showing the importance of including non-taxing states as a control. Column (3) shows the tax flight effect increases as the property tax rate increases: the $Tax\,State \times PreTaxTime \times Tax\ Rate$ parameter shows that the odds of being at the home airport is multiplied by 0.73 for a one unit increase in the tax rate. The other terms involving the tax rate are small and not statistically significant, indicating tax rates do not shape the propensity to be at the home airport in other time periods or in non-taxing states. Column (4) shows the results are robust to the inclusion of week fixed effects.

Tables 5 and 6 present estimates of the propensity of planes owners to engage in inter-state flight out-of or back-to their home airport. These estimates involve a restricted sample since the plane must be located at the home airport at the start of the week in the first case or in another state in the second case. The parameters in these tables are also consistent with tax flights. Column (1) shows that inter-state flights patterns are roughly comparable between taxing and non-taxing states. Column (2) shows that plane owners with a taxing home airport tend to fly out of their home airport to another state just before the assessment date ($Tax\ State \times PreTaxTime$ in Table 5has odds ratio of 1.42) and back to their home airport from another state just after the assessment date ($Tax State \times PostTaxTime$ in Table 6has odds ratio of 1.34, where $PostTaxTime$ which is defined analogously as $PreTaxTime$ except it is for the week following an assessment date). Column (3) shows that higher tax rates accentuate these effects, and column (4) shows that that the effects are robust to controls for week fixed effects.

Table 7 is the most direct measure of tax flights. Among planes which are located at their home airport in the beginning of the week, it considers whether the owner flies to another state in the current week and then returns in the following week. The dependent variable is an indicator for such roundtrip flights, and the sample is again only planes located at their home airport at the start of the week. The parameter on the $Tax\ State \times PreTaxTime$ interaction Column (1) shows such roundtrips are significantly more likely to occur in a taxing state just before the assessment date, relative to other times and to non-taxing states. Column (2) is a first step at looking at whether tax flights vary by the assessed value or by flight costs. $Tax\ State \times PreTaxTime$ is interacted with three categories of engine type which proxy for plane value. In order of increasing value, the categories are $Reciprocating/Piston$ plus $2/4 - cycle$ (the omitted category), $TurboProp$, $TurboFan$. Interstate roundtrips from the home airport around the assessment date are far more likely for more valuable planes, with the odds of a turbo fan plane taking such flight being about three times to that piston or n-cycle planes. The last row shows that the flights are also responsive to direct financial cost, namely the cost of fuel (see the Data Appendix for the construction of this variable). Column (3) shows these effects are robust to controlling for week fixed effects.

4.3 Validation of Estimates

While the previous section shows that plane owners behave in a manner consistent, with tax flights, there is no direct evidence that they are in face evading taxes. This sub-section addresses this point in two ways. First, I see whether planes with a home airport in the Kansas City metropolitan area are in fact missing from the tax roll. Second, I look at various placebo tests such as whether planes exempt from taxes also engage in tax flights. Finally, I consider the impact of a natural experiment which effectively created a new class of exempt planes in West Virginia towards the end of my sample. At this point, there are only results are preliminary numbers for the first approach.

The Kansas City metropolitan area includes seven counties in both Missouri and Kansas (both of these states tax planes). For each of these counties, I obtained the plane tax roll for each year over 2004-2009 (see the Data Appendix). These rolls include the owner name, address, plane and tax amounts for any plane on which property tax has been paid (one Kansas county has no airports and so has no tax roll). For each of these counties, I assemble the list of planes which have a home airport located in the county. I then see how many of these planes engage in a tax flight as defined in Table 7 and whether they are on the tax roll. About a fifth of the planes engage in tax flights, flying to another state from their home airport just before the assessment date and returning shortly thereafter. Strikingly, only about ten percent of these planes are listed on the tax rolls (about two thirds of all planes are on the tax roll). This is suggestive evidence that the flights are being used to evade taxes.

Remainder of this Section to be completed

4.4 Covariates of Tax Evasion

This Section to be completed

5 Conclusion

The preliminary evidence in this paper suggests that tax flights are a real and economically meaningful phenomenon. In the next revision of this work I hope to provide more specific evidence. In particular I will be making the following additions:

- **more formal estimates of the magnitude of tax flights:** after completing the assembly of the remaining data (local tax rates, approximating each airplane's assessed value, and geocoding the airplane locations), it will be possible to to estimate the specifications listed in the text. In particular this will allow me to estimate the increased likelihood of flights during the period just before and after the assessment date for a variety of different cases: taxed versus untaxed planes; high valuation versus low valuation planes; personal-owned taxed planes versus business-owned taxed planes.

- **weather (additional variation in flight patterns):** bad weather such as icy precipitation can force pilots to scrap planned trips, an important possibility around the most common assessment date of 1 January. While these conditions can typically be avoided using weather forecasts, sometimes fronts arrive more quickly or slowly than anticipated. I am in the process of assembling a database of actual weather as well as forecasts (three and seven days ahead) from NOAA at the airport-level.

- **plane operating costs (additional variation):** the paper includes one measure of flying cost, the cost of fuel which varies across both time and space. I will also add variation is in operating cost by plane model. This also should influence tax flights though this might be endogenous (strategic owners select into less costly planes) and much of the variation will be subsumed by the fixed effects.

- **tax evasion variation by demographic groups**: the data contain both the location of the plane and the address of the owner. While detailed background on the owner is unavailable, I can use fine grained geographic proxies (zip codes or census blocks) to proxy for demographics such as income. Variation in the benefits of tax flights can identify tax evasion for specific demographic groups. For example, tax evasion by income can be estimated from variation in tax rates and plane type (both of which impact the incentive to make a tax flight and do not covary one-for-one with local demographics). Understanding how evasion varies by group can help inform models of tax evasion and policy debates regarding tax rates.

- **natural experiment**: During the observation period West Virginia effectively exempted corporate-owned planes but did not change the taxation of those with other owners. There should be discontinuity in flight behavior around the assessment date for those planes affected by the change.
- **placebo tests**: I will also employ various placebo tests to help establish causality. Many states exempt certain planes from taxation, such as those owned by individuals or those owned by firms. Property taxes should have no impact on their flight patterns.

These revisions should provide a more precise measure of tax flights. Still there are other strategies which might be used to evade property taxes on airplanes. Owners might strategically hangar their planes in a non-taxing state, an attractive option for those who live near state borders (for example, owners in St. Louis may base planes in Illinois). Another possibility is that owners could put their airplane on the blocked list, which would prevent third parties including tax officials from monitoring their flight patterns. While this list has been private, the FAA for a short time made this list public (and at least subsets have been released under Freedom of Information Act requests). Exploring these and other tax evasion strategies are interesting topics for future work.

References

[1] Aircraft Bluebook Historical Value Reference (2010). Penton Media, Inc. http://www.aircraftbluebook.com

[2] Andreoni, James, Brian Erad, and Jonathan Feinstein (1998). “Tax Compliance.” *Journal of Economic Literature*. 36: 818-860.

[3] Artavanis, Nikolaos, Adair Morse, and Magarita Tsoutsoura (2011). “A New Method to Estimate Tax Evasion Using Financial Institution Lending: The Case of Greece.” Booth School working paper.

[4] Borenstein, Severin and Nancy Rose (1994). “Competition and Price Dispersion in the U.S. Airline Industry.” *Journal of Political Economy*. 102: 653-683.

[5] Brueckner, Jan (2002). “Airport Congestion When Carriers Have Market Power.” *American Economic Review*. 92: 1357-1375.

[6] Census (various years). *TIGER/Line Shapefiles*. http://www.census.gov/geo/www/tiger/index.html

[7] Chetty, Raj, Adam Looney and Kory Kroft (2009). “Salience and Taxation: Theory and Evidence.” *American Economic Review*. 99: 1145-1177.

[8] Chetty, Raj, John Friedman, Tore Olsen, Luigi Pistaferri (2011). “Adjustment Costs, Firm Responses, and Micro vs Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *Quarterly Journal of Economics*. 126: 749-804.

[9] ESRI ArcGIS (various years). *ESRI StreetMap Premium- North America*. http://www.esri.com.

[10] FAA (various years). *Aircraft Registry*. Personal correspondence.

[11] FAA (2009). *Aircraft Situation Display to Industry (ASDI)*. http://www.fly.faa.gov/ASDI/asdi.html

[12] FAA (2010). *Form 5010*: *Airport Master Record*. http://www.faa.gov/airports/airport_safety/airportdata_5010

[13] Fisman, Raymond and Shang-Jin Wei (2004). "Tax Rates and Tax Evasion: Evidence from 'Missing Imports' in China." *Journal of Political Economy*. 112: 471-496.

[14] Forbes, Silke and Mara Lederman (2009). "Adaptation and Vertical Integration in the Airline Industry." *American Economic Review*. 99: 1831–1849.

[15] Goolsbee, Austan and Chad Syverson (2008). "How Do Incumbents Respond to the Threat of Entry? The Case of Major Airlines." *Quarterly Journal of Economics*. 123: 1611-1633.

[16] Kleven, Henrik, Martin Knudsen, Claus Kreiner, Soren Pedersen, and Emmanuel Saez (2011). "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark." *Econometrica*. 79: 651-692.

[17] Kleven, Henrik and Mazhar Waseem (2011). "Tax Notches in Pakistan: Tax Evasion, Real Responses, and Income Shifting." LSE working paper.

[18] Lincoln Institute (2010). *Significant Features of Property Tax*. George Washington Institute of Public Policy. http://www.lincolninst.edu/subcenters/significant-features-property-tax/Report_TaxRates.aspx

[19] Maponics (2010). *ZIP4 Database Premium plus Centroids*. http://www.maponics.com.

[20] Marion, Justin and Erich Muehlegger (2008). "Measuring Illegal Activity and the Effects of Regulatory Innovation: Tax Evasion and the Dyeing of Untaxed Diesel." *Journal of Political Economy*. 116: 633-666.

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

[22] Merriman, David (2010). "The Micro-Geography of Tax Avoidance: Evidence from Littered Cigarette Packs in Chicago." *American Economic Journal: Economic Policy*. 2: 61-84.

[23] NBAA (2010). *Block Aircraft Registration Request (BARR) Program.* http://www.nbaa.org/ops/security/barr/

[24] Slemrod, Joel (2007). "Cheating Ourselves: The Economics of Tax Evasion." *Journal of Economic Perspectives*. 21: 25-48.

[25] Slemrod, Joel and Shlomo Yitzhaki (2002). "Tax avoidance, evasion, and administration," in: A. J. Auerbach & M. Feldstein (ed.), *Handbook of Public Economics*, edition 1, volume 3, chapter 22, Elsevier, 1423-1470.

[26] USPS (2010). *Unmatched ZIP+4 Data File.* Personal correspondence with Maponics, LLC.

[27] Winston, Clifford and Steven Morrison (1995). The Evolution of the Airline Industry. Washington D.C.: Brookings Institution Press.