How Do Airlines Cut Fuel Usage, Reducing their Carbon Emissions?

by

Jan K. Brueckner
University of California, Irvine
jkbrueck@uci.edu

Matthew E. Kahn
University of Southern California
kahnme@usc.edu

Jerry Nickelsburg
UCLA
jerry.nickelsburg@anderson.ucla.edu

May 2023, revised April 2024

Abstract

Airline fuel consumption is costly for the firms and for society as well due to a climate-change externality. We study how fuel-price changes affect cost-minimizing choices by airlines that have implications for the extent of this externality. The airline industry's capital stock can be easily inventoried as a set of long-lived, durable aircraft. This portfolio approach allows us to study the utilization and composition of the capital stock at a highly disaggregated level. Changes in airline operations directed toward conserving fuel can be an important path toward lower emissions.

How Do Airlines Cut Fuel Usage, Reducing their Carbon Emissions?

by

Jan K. Brueckner, Matthew E. Kahn, and Jerry Nickelsburg†

1. Introduction

Jet fuel is a major expense for commercial airlines. In 2019, the U.S civilian fleet consumed 12.2 billion gallons for domestic flights, representing a total expenditure of $24.3 billion.[1] Since this expense accounts for approximately 25% of operating costs, airlines have strong incentives to manage their fuel usage. However, carriers do not internalize a crucial externality generated by combustion of jet fuel: the contribution of the resulting greenhouse gas emissions (mainly CO_2) to global warming and climate change. These emissions are directly proportional to an aircraft's fuel consumption. While airlines contribute only a bit more than 2% of GHG emissions in the US (representing 8% of transportation emissions),[2] their prominence in the public eye draws attention to this contribution and its harm. Moreover, as the usage of electric automobiles grows, reducing total emissions from the transportation sector, the contribution of aviation will become more prominent.

Biofuels offer a possible path toward lower airline emissions, but their high cost makes this solution currently impractical. However, the Biden administration, as part of its broader efforts to decarbonize the transportation sector, is subsidizing the development of sustainable aviation fuels (SAF). Conceivably, SAF will become economical for airline use by mid-century.

Pricing of emissions is an alternative. While this approach is unlikely to be adopted in the US, it is followed on a large scale in Europe, where intra-EU flights are subject to the EU's Emissions Trading System. Independently, ICAO (a unit of the United Nations) launched a worldwide carbon-offset program for airlines called CORSIA, where airlines purchase offsets for

† This paper builds on the earlier work of Kahn and Nickelsburg (2016) through use of an additional five years of data and new analysis. We thank Kangoh Lee, Joshua Graff Zivin and several referees for helpful comments, but the usual disclaimer applies.

[1] https://www.transtats.bts.gov/fuel.asp

[2] See https://www3.epa.gov/otaq/documents/aviation/420f15023.pdf.

emissions above a 2020 baseline value.[3] The program is voluntary until 2026 but mandatory thereafter.

While awaiting the emergence of affordable SAF, improvements in aircraft fuel efficiency offer the most effective current path to lower airline emissions. For example, the Airbus A320, which began service in the late 1980s, emits 46% less CO_2 than a Boeing 727-200, a model that retired from service long ago. The Airbus A320neo, an updated version of the A320 introduced recently, has more fuel-efficient engines and generates 18% less emissions than the earlier model. The Boeing 737 MAX offers a similar fuel-efficiency improvement over previous 737 models, and several newer widebody aircraft yield analogous gains (see below for more detailed information).

Since aircraft themselves thus appear to be the best sources of near-term improvements in airline emissions, it is important to understand how carriers adjust aircraft utilization and the composition of their fleets in response to fuel price dynamics. Their privately optimal choices show how airline emission externalities can shrink even in the absence of Pigouvian taxation. By estimating the scope of these effects, we can then better gauge how airlines would respond to policy measures such as an increase in fuel taxes. We return to this question in the paper's concluding section.

While all corporations are major producers of greenhouse gas emissions, data constraints usually limit our ability to explore at a detailed level the possible channels for industrial pollution reduction. The US Census of Manufacturers surveys firms on their annual energy consumption, but the survey instrument does not allow researchers to explore choices at the intensive or extensive margins that together determine aggregate energy consumption.[4] In contrast, available data allow the airline industry's capital stock to be easily inventoried as a discrete set of long-lived, durable aircraft. This portfolio approach allows us to study the composition and utilization of the capital stock at a highly disaggregated level.

[3] CORSIA stands for Carbon Offsetting and Reduction Scheme for International Aviation. Although structured differently, this plan is equivalent to requiring the purchase of allowances under an ETS-style system. See `http://www.icao.int/environmental-protection/CORSIA/Pages/default.aspx`

[4] Data at the six digit NAICS/year level can be used to calculate energy efficiency gains over time and to compare energy efficiency across industries. See `https://www.nber.org/research/data/nber-ces-manufacturing-industry-database`.

Changes in airline operations directed toward conserving fuel can be an important path toward lower emissions, and this channel is a main focus of the present paper. Previous evidence of such conservation effects is given by Brueckner and Abreu (2017, 2020) and Fukui and Miyoshi (2017), who show that airline fuel usage falls as the fuel price rises. Brueckner and Abreu (2017) find this effect at the airline level, holding miles flown and fleet characteristics (and thus fuel efficiency) constant, while Brueckner and Abreu (2020) find the same effect at the aircraft-model level. Both results, along with those of Fukui and Miyoshi (2017), provide indirect evidence of conservation efforts.[5] These efforts, which are not measured directly, can include lower flight speeds, taxiing on one engine, carrying less (heavy) reserve fuel, installation of fuel-saving winglets, and favoring the more fuel-efficient aircraft in the airline's fleet.[6]

One purpose of the current paper is to provide direct, rather than indirect, evidence of airline fuel conservation in response to higher fuel prices, using data from the 1991-2019 period. Central to our exercise is the recognition that the cost impact of a higher fuel price will depend on the fuel efficiency of individual aircraft. Accordingly, for each year, we compute gallons used per seat-mile (gallons PSM) for each airline/aircraft-type combination, and then multiply this value by the current real fuel price per gallon. The result is the fuel cost per seat-mile (fuel cost PSM) by aircraft type and airline. This measure is used as an explanatory variable in several regressions that focus on particular aspects of airline operations, providing evidence of fuel-conservation efforts.

Fast freeway drivers know that they can conserve fuel by driving slower, a gain that is larger when the car's overall fuel efficiency is low. Our first regression investigates a related factor in airline operations. It asks whether aircraft with higher fuel cost PSM are flown at lower speeds to conserve fuel.[7] Since the regression's explanatory variable, fuel cost PSM,

[5] While research on fuel economy impacts for airlines is scarce, a bigger literature focuses on the private automobile fleet and the public bus fleet. See Knittel (2012) and Li, Kahn and Nickelsburg (2015).

[6] Fageda and Texeido (2022) investigate the effects of the EU's Emissions Trading System on airline emissions. Using a difference-in-difference approach, they show that emissions fell after 2013 on intra-EU routes, which had then become subject to the ETS, relative to emissions on routes with one endpoint outside the EU, which were exempt. They find that most of the decrease came from a reduction in intra-EU traffic in response to the pricing of emissions.

[7] Aircraft fuel consumption as a function of speed takes a parabolic form, as seen in Aktürk, Atamtürk and Gürel (2014) and Matsuno and Andreeva-Mori (2020), with consumption rising beyond the Maximum Range Cruise speed (MRC). See also Boeing (2017) as well as Moskwa (2008) for media coverage of aircraft speeds.

depends on both fuel efficiency (gallons PSM) and the current fuel price, both these elements contribute to the expected effect on speed. The results show the expected negative relationship between speed and lagged fuel cost PSM, providing evidence that airlines limit flight speeds as a way of conserving fuel.[8]

Because gallons PSM will itself partly depend on speed, a "semilag" of fuel cost PSM is used in the regression to avoid reverse causality from speed to the gallons PSM component of fuel cost PSM. This variable is generated by multiplying the current fuel price by the one-year lag of gallons PSM, the endogenous component of fuel cost PSM. In addition, since the regression uses aircraft-type, airline, and year fixed effects, the estimated negative effect holds the aircraft type constant, being generated by variation in fuel cost PSM across years and airlines within aircraft types. A recent paper by de Almeida and Oliviera (2023) carries out a related empirical inquiry using Brazilian data.[9]

Since a lower flying speed will reduce the number of flights an aircraft can operate each period, a high fuel cost PSM is expected, via lower speeds, to reduce aircraft utilization. Utilization could also be reduced by operating an aircraft fewer hours per period in response to a high fuel cost PSM. In other words, aircraft with high costs would spend more time on the ground than their more fuel-efficient counterparts. To test for utilization effects through these two channels, the second regression relates annual available seat-miles for an aircraft type to its semi-lagged fuel cost PSM, finding the expected negative relationship. Thus, when fuel cost PSM is high due to some combination of low fuel efficiency and high fuel prices, airlines conserve fuel usage through lower aircraft utilization. Like the first regression, this one uses aircraft-type, airline, and year fixed effects, so that the negative utilization effect again holds aircraft type constant. Both the speed and utilization regressions are motivated by a

[8] This finding matches anecdotal evidence provided by a colleague at the National University of Singapore, whose friend is a pilot for MYAirline and previously AirAsia, both low-cost carriers. The friend told him that "it is routine for the airline to monitor fuel prices and instruct pilots to fly at lower cruise speeds when prices are high, to reduce the amount of fuel used..."

[9] Their study focuses on the determinants of aircraft speed. The regressions measure speed in two alternate ways: the planned speed given in the aircraft's flight plan, and the actual speed computed as the ratio of flight time to distance. While both speeds are higher when the aircraft is a new fuel-efficient type (mirroring our results), the fuel price only has the expected negative effect on the actual speed, not on the planned speed (which is more likely to reflect airline conservation decisions).

theoretical model presented in section 2 of the paper.

In addition to presenting these results on fuel conservation, the paper explores another channel by which fuel prices can reduce emissions: replacement of older, fuel-inefficient aircraft with new planes. We use two approaches in analyzing fleet replacement. First, we attempt to measure the effect of fuel prices on the ages and fuel efficiencies of aircraft in an airline's fleet, analysis that extends earlier work by Goolsbee (2008) on the retirement of the Boeing 707. Replacement is alternately captured by (*i*) the annual change in an airline fleet's average gallons PSM and (*ii*) the annual change in an airline fleet's average aircraft age. As older aircraft are replaced by newer, more fuel-efficient planes, both changes are negative. The regressions relate these variables to the annual change in the real fuel price as well as the lagged change. Note that, in contrast to the speed and utilization regressions and the ones described next, these regressions are carried out at the airline/year level rather than at the aircraft-type/airline/year level.

The second approach focuses on the rate of drawdown of older aircraft types, as well as the rate of buildup of new types. In the drawdown regression, the dependent variable is the percentage annual drop in the count of an older aircraft type in an airline's fleet when the count is falling. One explanatory variable is "relative gallons," equal to gallons PSM for that type divided by average gallons PSM in the airline's fleet. The other main explanatory variables are the fuel price and the interaction of the fuel price and relative gallons. The buildup regression is the mirror image of the drawdown regression, focusing on aircraft types whose count is rising.

The paper's final contribution is a presentation of descriptive evidence tracking the fates of aircraft once they are retired from major airline fleets. These fates include transfer to other airlines around the world or scrappage, which often provides a source of parts for aircraft remaining in a fleet. While banning the production of polluting aircraft is obviously not an option, in contrast to plans in California and the UK to ban production of non-electric vehicles (Holland, Mansur and Yates (2021)), retirement and scrappage of old inefficient planes can help limit aviation's climate impact. Policies for accelerating the aircraft scrappage documented in our analysis are worthy of future study.

The data for the flight speed and fleet utilization and regressions are derived from the T2 database of the US Bureau of Transportation Statistics, which shows annual fuel usage, flight hours, and flight distances by aircraft type and airline.[10] For the replacement regressions, these data are supplemented by annual, hand-collected data on aircraft counts and average ages by type for each airline, drawn from non-government sources described below (this time-intensive data effort is itself a major contribution of the paper). The sample consists of data on 17 major airlines over the 1991-2019 period.

The plan of the paper is as follows. Section 2 presents a theoretical model, while section 3 discusses the data sources and variable definitions. Section 4 presents descriptive statistics, and section 5 presents the regression results. Section 6 discusses the fates of retired aircraft, and section 7 offers conclusions.

2. Theoretical model

This section presents a theoretical model that motivates the empirical analysis of aircraft speed and utilization. The model does not treat aircraft replacement, our second empirical focus. Suppose that an airline wishes to operate F total flights per period using two aircraft types, with type 1 being more fuel efficient than type 2. The airline owns N_1 aircraft of type 1 and N_2 aircraft of type 2, and both types have the same number of seats. Distance is the same for all flights. The flight speeds of the two aircraft types are denoted v_1 and v_2, and they are choice variables of the airline.[11] The fuel cost for type-i aircraft is denoted $c_i(v_i)$, with $c_i' > 0$ and $c_i'' > 0$, indicating that costs rise at an increasing rate as speed increases. Suppose that the functions c_1 and c_2 differ only by a multiplicative factor, so that $c_i(v_i) = \beta_i c(v),\ i = 1, 2$, where $c', c'' > 0$ and $\beta_1 < \beta_2$ (type 1 is more fuel efficient).

A lower flight speed reduces the number of flights that an aircraft can operate per period. Let T denote the length of a period, which is best viewed as a month or a year. With T measured in hours and D denoting the common flight distance, flights per period for an

[10] See https://www.transtats.bts.gov/Fields.asp?gnoyr_VQ=FIH. These data are available quarterly, but a quarterly focus seems unnecessary (and perhaps generates uninformative noise) since the data cover almost a 20-year period.

[11] Speed differs across the cruise, takeoff and landing portions of a flight, with these variables representing average speeds.

aircraft equals

$$f(v) \quad = \quad T \div \text{hours/flight} \quad = \quad T \div \frac{\text{miles/flight}}{\text{miles/hour}} \quad = \quad T \div (D/v) \quad = \quad (T/D)v \quad \equiv \quad \alpha v, \tag{1}$$

where $\alpha = T/D$. Thus, flights per period is proportional to aircraft speed. Using all this information, total fuel cost for an airline equals $N_1 f(v_1)\beta_1 c(v_1) + N_2 f(v_2)\beta_2 c(v_2)$, or the sum across aircraft types of the number aircraft × flights per aircraft × fuel cost per flight (with the f terms given by (1)).

Revenue per flight is denoted R, and it is assumed to be independent of speed. While a dramatic speed reduction would noticeably lengthen flight duration, reducing consumer willingness-to-pay, the effect of smaller fuel-conserving reductions are likely to be imperceptible to consumers, justifying the fixed-R assumption. The airline's total revenue is then fixed at RF, where F is again the fixed flight total. Ignoring non-fuel costs, the Lagrangean expression for the airline's profit maximization problem is

$$RF \; - \; [N_1 f(v_1)\beta_1 c(v_1) \; + \; N_2 f(v_2)\beta_2 c(v_2)] \; + \; \lambda[N_1 f(v_1) \; + \; N_2 f(v_2) \; - \; F], \tag{2}$$

where the second expression is total fuel cost and where λ is Lagrange multiplier, which multiplies the expression embodying the total flight constraint (which is set at zero).

The first-order conditions for choice of v_1 and v_2 are

$$N_i[f'(v_i)\beta_i c(v_i) \; + \; f(v_i)\beta_i c'(v_i) \; - \; \lambda f'(v_i)] \quad = \quad 0, \quad i = 1, 2. \tag{3}$$

Substituting for f and $f' = \alpha$, (3) becomes

$$\alpha\beta_i c(v_i) \; + \; \alpha v_i \beta_i c'(v_i) \quad = \quad \lambda\alpha, \quad i = 1, 2. \tag{4}$$

Dividing through by α, and then dividing the equation for $i = 1$ by the equation for $i = 2$, (4) can be written, after extracting the β's, as

$$\frac{c(v_1) \; + \; v_1 c'(v_1)}{c(v_2) \; + \; v_2 c'(v_2)} \quad = \quad \frac{\beta_2}{\beta_1} \quad > \quad 1. \tag{5}$$

Since $c(v) + vc'(v)$ is increasing in v given $c', c'' > 0$, satisfaction of (5) requires $v_1 > v_2$. Therefore, the less fuel-efficient aircraft type (type 2) is flown slower than type 1.

Type 2's lower speed translates into fewer flights per period, with $f(v_2) = \alpha v_2 < \alpha v_1$. But it is possible that the airline further reduces utilization of type-2 aircraft by operating them less intensively otherwise. This channel can be captured by letting $A_i \leq N_i$ denote the effective number of aircraft of type i operated by the airline. For example, if type-i planes are operated for only half of their feasible hours, then A_i would equal $N_i/2$.

To capture this other utilization channel, the maximization problem in (2) can be recast by replacing N_i by A_i and adding the constraints $N_i \geq A_i,\ i = 1, 2$, with Lagrange multipliers $\rho_i \geq 0,\ i = 1, 2$. A_1 and A_2 then become choice variables, and their first-order conditions are

$$f(v_i)(\lambda \ - \ \beta_i c(v_i)) \ \ = \ \ \rho_i, \quad i = 1, 2. \tag{6}$$

To derive the implications of (6), suppose that $\beta_1 c(v_1) < \beta_2 c(v_2)$ holds, which says that fuel cost per flight is lower for type-1 aircraft. When $v_1 > v_2$, this relationship is not guaranteed to hold, but the outcome seems natural given higher type-1 fuel efficiency ($\beta_1 < \beta_2$). Then, $\lambda - \beta_1 c(v_1) > \lambda - \beta_2 c(v_2)$ holds in (6), and this inequality in turn implies that ρ_i cannot be zero for both aircraft types, with at least one type then fully utilized. Type 1 (type 2) will be fully (partially) utilized, when $\lambda - \beta_1 c(v_1) > 0$ and $\lambda - \beta_2 c(v_2) = 0$ hold, implying $\rho_1 > 0$ ($A_1 = N_1$) and $\rho_2 = 0$ ($A_2 < N_2$). Therefore, beyond a negative utilization effect due to lower speed, the low-efficiency aircraft type may not be flown as much as possible, spending more time on the ground than its type-1 counterpart. While it would appear that this outcome is less likely when the total flight target F is high, it seems possible when there is more slack in the airline's optimization problem.

The possibility of corner solutions must be considered in this modified optimization problem. One uninteresting case is where the stock of efficient type-1 aircraft is large enough to meet the flight target with no use of type-2 aircraft, in which case $A_2 = 0$. However, a different corner solution is certain to arise in an unexpected place: flight speed. When $\beta_2 c(v_2) - \lambda = 0$ is substituted into (4), the speed first-order condition, the condition cannot hold as an equality, with the implication that v_2 should be as small as possible. This outcome implies the need for a minimum-speed constraint $v_i \geq \overline{v},\ i = 1, 2$, which will be binding for type-2 aircraft. With the model now more complex, $v_1 > \overline{v}$ (yielding $v_1 > v_2$) cannot be established unambiguously, but this outcome again seems natural.

3. Data and Variable Definitions

To compute aircraft speed, fuel efficiency, and utilization, we use data from the T2 data set of the U.S. Bureau of Transportation Statistics (BTS).[12] For each year, airline, and aircraft type, this source gives fuel usage, revenue aircraft miles flown, revenue aircraft hours airborne, and available seat-miles, yielding 2,058 observations. Letting *fuel_price* denote annual average aviation fuel price in constant dollars per gallon,[13] the following additional variables are computed using the BTS information:

$$\begin{aligned}
speed &= \frac{revenue\ aircraft\ miles\ flown}{revenue\ aircraft\ hours\ airborne} \\
avl_seat_miles &= available\ seat\ miles \\
gallons_seat_mile &= \frac{fuel\ usage}{avl_seat_miles} \\
cost_seat_mile &= gallons_seat_mile \times fuel_price
\end{aligned} \tag{7}$$

Again, we generate each of these variables by aircraft type (a), airline (c, for carrier), and year (t), although these subscripts in (7) are suppressed for readability. To reduce measurement error, observations with values of *speed* and *cost_seat_mile* in the top and bottom 1% of their respective ranges are deleted. Note that the speed variable defined above is ground speed, as opposed to air speed (head winds make air speed higher than ground speed).

An alternative measure of an aircraft's operating cost would be cost per passenger mile (*cost_pax_mile*), which is generated by replacing *avl_seat_miles* in the third line of (7) with

[12] See `https://www.transtats.bts.gov/Fields.asp?gnoyr_VQ=FIH`. These data are available quarterly, but a quarterly focus seems unnecessary (and perhaps generates uninformative noise) since the data cover almost a 20-year period.

[13] The real fuel price is the average of monthly fuel prices during each year from `https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=EMA_EPJK_PWG_NUS_DPG&f=M`, adjusted by the annual consumer price index.

passenger miles (*pax_miles*), which counts occupied seats. The result (*gallons_pax_mile*) is then used in place of *gallons_seat_mile* in the fourth line of (7). Since *pax_miles* = $L \times$ *avl_seat_miles*, where L is the load factor (percent of seats filled), it follows that *cost_pax_mile* = $(1/L) \times$ *cost_seat_mile*. Because of *cost_pax_mile*'s explicit dependence on load factor, we prefer *cost_seat_mile* as a cost measure (either way, the regression results are similar). However, it should be noted that both measures *implicitly* depend on load factor, because fuel usage in both numerators will be higher when the aircraft is more fully loaded. Interestingly, taking both the implicit and explicit dependencies into account, L is likely to affect the two measures in opposite directions.[14]

It should be noted that an inherent limitation in our empirical analysis is that the fuel price faced by an airline may differ from the current price because of fuel hedging, which protects carriers against unfavorable price movements. While systematic data on hedging by individual carriers over our entire sample period is not available, Merkert and Swidan (2019) provide evidence for 2019, which shows that American, Delta and United were unhedged in this year (although the discussion suggests they were hedged earlier) while Southwest's fuel costs were 64% hedged in 2019. Hedging is a short-term activity, not extending far into the future, and as shown by the Southwest case, airlines tend to only hedge a portion of their fuel costs. Unmeasured hedging may be partly picked up through our carrier fixed effects, although variation in hedging over time will not be adequately captured. The upshot is that fuel-cost hedging is a source of measurement error in our analysis.

Because the online BTS data are incomplete, we use non-government data sources to capture two additional pieces of information for each aircraft type in an airline's fleet: for each year, the count (number of planes) for that type and the average age of the planes of each type. Hand collection of these data, which was extremely time intensive, relied on three

[14] To see this point, suppose that an aircraft's fuel usage, holding speed constant, can be written as $\eta S^{\gamma} L^{\nu}$ where S is seat-miles, η is an index of fuel (in)efficiency, and γ and ν are less than 1 (the empirical results of Brueckner and Abreu (2017, 2020) justify these magnitudes). Using this expression, *gallons_seat_mile* = $\eta S^{\gamma} L^{\nu}/S = \eta S^{\gamma-1} L^{\nu}$. On the other hand, *gallons_pax_mile* = $\eta S^{\gamma} L^{\nu}/(LS) = \eta S^{\gamma-1} L^{\nu-1}$, using *pax_miles* = LS. Therefore, after multiplying by the fuel price and recalling $\nu < 1$, the load factor L has opposing effects on the two cost measures.

sources: Planespotters.net, Planelist.net, and Airfleets.net.[15] Planespotters provides a list of aircraft types at each airline, with an introduction-to-service date and, for many of the aircraft, a removal-from-service date. Planelist.net provides a data check on the Planespotters data. It also traces each aircraft by manufacturer's line number through its entire life, including ownership and usage. The data in Planelist.net were the default in the event of discrepancies between the Planespotters and Planelist sources. Those discrepancies were only in the usage of the aircraft after removal from service and in the removal-from-service date. Airfleets.net provided backup data on the fleet sizes and a final check on the veracity of the data. Compilation of the aircraft count and average age data is by itself a major contribution of the paper.

Aircraft were entered into a type's count if they were in the fleet for more than six months in a year. If entry occurred after June, aircraft were counted as entering the fleet in the following year. The same rule was used for aircraft exits.[16] Moreover, the entry date was used to determine the effective age of the aircraft rather than relying on the calendar age from completion of manufacturing. In addition, aircraft acquired through a merger or purchase of another carrier that were removed from service within a year of the acquisition were not counted as being part of the acquiring carrier's fleet. Aircraft that entered service directly from the manufacturer in the first six months of the year were given an age of 1/2 year for the first year.

The data on aircraft counts and age by type were used to compute variables for the fleet replacement regressions. Introducing subscripts, let $tot\_count_{ct} = \sum_a count_{act}$ denote the total count of planes across all aircraft types a in airline c's fleet in year t, where $count_{act}$ is the count of aircraft type a for the airline. Letting age_{act} denote the age of the airline's type-a aircraft, the average age of aircraft in a carrier's fleet in year t is given by[17]

$$avg\_age_{ct} = \frac{\sum_a count_{act} \times age_{act}}{tot\_count_{ct}} \tag{8}$$

[15] These URLs lead directly to the sources, although the name of the Planelist website is Airlinelist.

[16] While data exist on the month of acquisition of an aircraft, that month does not necessarily correspond to the entry (beginning-of-service) date due to pilot training and marketing considerations.

[17] Note that age_{act} is itself an average, since planes of a given type may have been produced in different years.

In addition, the average gallons PSM of aircraft in an airline's fleet is given by

$$avg\_gallons\_seat\_mile_{ct} \quad = \quad \frac{\sum_a count_{act} \times gallons\_seat\_mile_{act}}{tot\_count_{ct}} \tag{8}$$

The drawdown and buildup regressions use the percentage changes of the aircraft-type count, as follows:

$$drawdown_{act} \quad = \quad \begin{cases} \frac{count_{act-1} - count_{act}}{count_{act-1}} & \text{if } count_{act-1} - count_{act} > 0 \\ 0 \quad otherwise \end{cases} \tag{9}$$

$$buildup_{act} \quad = \quad \begin{cases} \frac{count_{act} - count_{act-1}}{count_{act-1}} & \text{if } count_{act} - count_{act-1} > 0 \text{ and } count_{act-1} > 0 \\ 0 \quad otherwise \end{cases} \tag{10}$$

Note that $drawdown_{act}$ and $buildup_{act}$ are defined to be positive and pertain to types whose counts are falling and rising, respectively. Observe also that, to avoid dividing by zero, $buildup_{act}$ is not computed for the initial aircraft of a type added to the fleet. An additional variable used in the drawdown and buildup regressions is the relative-gallons measure:

$$rel\_gallons_{act} \quad = \quad \frac{gallons\_seat\_mile_{act}}{avg\_gallons\_seat\_mile_{ct}} \tag{11}$$

4. Descriptive statistics

Table 1 provides summary statistics for most of these variables as well as for dummy variables for the 17 airlines. Note that the sample size for the variables *fuel_price*, *avg_age* and *avg_gallons_seat_mile* is smaller because they vary only by airline and year, not by aircraft type, airline and year. Observe also that the maximum aircraft speed in the sample is just below 540 miles per hour, a value achieved by United 747-400 aircraft in 2011. This value is close to the 580 mph cruising speed of the aircraft, an outcome that is possible because its long flight distances reduce the importance of the slower takeoff and landing phases. With most aircraft flying shorter distances, these slower phases comprise a greater share of the flight distance, leading to a lower average speed of 455 mph across the entire sample.

Table 2 shows the frequencies of the different aircraft types in the sample (the number of carrier/year appearances) along with average gallons per seat-mile for each type. As can be seen, vintage Boeing and Douglas narrow-body aircraft (Boeing 727, 737-100/200 and DC 9-10/30/40/50) have gallons PSM in the 0.023-0.035 range. Later Boeing 737 models (the -300/-400/-500/-700 variants) have gallons PSM somewhat below that range (0.015-0.019). McDonnell-Douglas successors to the DC-9 (MD-80/81/82/83/88) also have gallons PSM in this range, as do contemporaneous Airbus narrow body aircraft (A318, A319, A320-100/200). The newest narrowbody planes from both manufacturers, the Boeing 737-800/900/Max 800/Max 900 and Airbus 320neo and 321neo models, are notably more fuel efficient than their predecessors, with gallons PSM in the 0.010-0.013 range. Earlier Boeing 757-200/300 models, like the A321, are relatively large narrowbody aircraft, and they had somewhat higher gallons PSM, in the 0.013-0.014 range.

Vintage Boeing widebody aircraft (747-100/200/300/SP) along with the more modern 747-400 version had relatively high gallons PSM, in the 0.017-0.025 range. Later Boeing models (767-200/300/400) and earliest Airbus widebody (A300) were more fuel efficient than the 747s, with values in the 0.014-0.018 range, while values for later widebody models (Boeing 777-200/300 and Airbus 330-100/200/300/333) were not much lower. The latest widebodies from these manufacturers (Boeing 787-800/900/10 and Airbus 330-900 and 350-900) are considerably more fuel efficient, with gallons PSM in the 0.011-0.015 range. The earlier, less-successful widebody aircraft (DC-10-10/30/40, MD-11 and Lockheed L1011) were relatively fuel inefficient, with gallons PSM in the 0.018-0.023 range.

Figures 1 and 2 illustrate the improvement in aircraft fuel efficiency over the 1991-2019 sample period. Figure 1 graphs average gallons per seat-mile for the three airlines that are currently the largest: American, Delta and United. As can be seen, for each carrier, average gallons per seat-mile fell from above 0.018 in 1991 to below 0.016 by 1991. Figure 2 provides more detail for American, showing the distribution of gallons per seat-mile across aircraft types for the years 1995, 2003, 2011 and 2019. As can be seen, the distributions shift to the left over time, indicating greater aircraft fuel efficiency.

Figure 3 shows time path of the real fuel price per gallon over the sample period. From a low of \$0.62 per gallon in 1998, the price rose to \$1.55 per gallon in 2012, then fell to \$0.86 per gallon in 2016 while rising somewhat thereafter, reaching \$0.98 per gallon at the end of the sample period.

5. Empirical models and regression results

5.1. Speed and available seat-miles regressions

As explained above, the specifications of these regressions are very simple. Letting z_{act} denote either speed or available seat-miles for aircraft type a operated by carrier c in year t,

$$lz_{act} \;=\; \theta\, lcost\_seat\_mile_{act-1} \;+\; \delta_a \;+\; \mu_c \;+\; \omega_t \;+\; u_{act}, \tag{12}$$

where the variables are used in log form, indicated by an "l" preceding the variable name, and where δ_a, μ_c, and ω_t are aircraft-type, carrier, and year fixed effects (u_{act} is the error term). Note that the lag of $lcost\_seat\_mile$ in (12) is actually a "semilag", as explained above, with the variable based on the *current* fuel price but a *lag* of gallons_seat_mile. In other words, $lcost\_seat\_mile_{act-1}$ equals $log(fuel\_price_t \times gallons\_seat\_mile_{act-1})$. Recall that the lag is needed because of possible reverse causation running from speed to $cost\_seat\_mile$. This effect creates negative correlation between $lcost\_seat\_mile$ and the regression error term, leading to downward bias in the variable's coefficient. In particular, a large value for the error term, which leads to high speed, will in turn reduce $gallons\_seat\_mile$, creating the negative correlation between $lcost\_seat\_mile$ and the error term. By using lagged $gallons\_seat\_mile$ instead of the current value along with the current fuel price to create a semi-lagged variable, this bias may be prevented. While another approach to dealing with this bias would rely on an instrumental variable, a usable instrument was not available.

It is important to note in (12) that aircraft-type fixed effects will help to control for differences in average speeds across different types of planes. Larger aircraft cruise somewhat faster, and their longer stage lengths mean a greater share of flight time is spent at these higher cruise speeds. Within aircraft types, by contrast, variation in speed will arise partly because of differences in $cost\_seat\_mile$ induced by fuel-price variation over the sample years, which is captured by the year fixed effects. It might then seem that all speed variation will be absorbed by the aircraft-type and year fixed effects, leaving no additional variation to be explained by *lcost_seat_mile*. But recall that *gallons_seat_mile* also depends on an aircraft's load factor. Therefore, load factors that vary by aircraft type and airline provide an independent force that affects fuel usage and hence speed throught *lcost_seat_mile*.[18]

Fuel consumption for a given aircraft type might also vary as a result of differences across airlines in route assignments and thus route distances. For example, one airline might use its 737-800s both for transcontinental flights as well as shorter flights, while another airline with 757-200s in its fleet might assign these planes to transcontinental flights while using its 737-800s only for shorter flights. With *gallons_seat_mile* lower for longer stage lengths, such differences might provide an additional source of variation in this variable within aircraft types across airlines. However, since a given plane is likely fly on routes with a variety of distances (within its range) over a single day, different aircraft of a given type probably fly routes with similar average lengths over a year. If so, the preceding argument loses force, and aircraft fixed effects will capture most of the route-distance variation across plane types.

Similarly, differences in route assignments across airlines could potentially expose planes of a given type to different weather and wind environments, with differences in the average velocity of head or tail winds leading to differences in *gallons_seat_mile*. While route assignment decisions might therefore generate an additional force beyond stage length that generates variation in *gallons_seat_mile* within aircraft types across airlines, the discussion above tends to undercut this possibility, with route variety exposing aircraft of a given type to similar average weather over a year.

The results of estimating (12) are shown in columns 1–3 of Table 3. Column 1 shows results from a speed regression using the current value of *lcost_seat_mile_lag*, which may lead to a downward-biased coefficient. The coefficient is negative and significant at the 1% level. Column 2 replaces this variable with its semi-lag, denoted *lcost_seat_mile_semilag*, and the coefficient is again significantly negative at the 1% level but smaller in absolute value, a difference consistent with the predicted direction of bias. Thus, a high cost PSM reduces

[18] Attempts to use load factor as an instrument for *lcost_seat_mile* were unsuccessful.

aircraft speed. The *lcost_seat_mile_semilag* coefficient of -0.125 indicates that a 10% increase in the semi-lagged cost PSM reduces flight speed by 1.25%. At an average speed of 455 mph, this reduction equals 5.7 mph.[19] Note that, because *cost_seat_mile* has a multiplicative form, this 10% increase could come either from a 10% increase in the fuel price or a 10% increase in (lagged) gallons PSM, yielding a speed comparison between a given aircraft and a less fuel-efficient plane.

Column 3 of Table 3 shows that, as predicted, aircraft with a high semi-lagged *lcost_seat_mile* are less utilized, generating fewer available seat-miles per year. The coefficient of -1.477 (again significant at the 1% level) shows a 10% increase in cost reduces *avl_seat_miles* by 15%. In this regression, use of aircraft-type fixed effects controls for innate variation in available seat-miles across types due to differences in stage lengths and ground times (factors that vary across long- and short-haul aircraft). Once again, the cost effect is further identified by variation in *cost_seat_mile* within aircraft types across airlines and years (caused by fuel-price, load factor and route and weather variation).

As seen in the theoretical discussion of section 2, aircraft with a higher cost per seat-mile may spend more time on the ground than lower-cost planes, not being operated as intensively. But the model also showed that available seat-miles are mechanically related to speed, given that a lower speed allows fewer flights per period. Therefore, in addition to an available-seat-miles regression with the form of (12), we report a regression with log speed as an additional covariate:

$$lavl\_seat\_miles \;=\; \theta\, lcost\_seat\_mile_{act-1} \;+\; \tau\, lspeed_{act} \;+\; \delta_a \;+\; \mu_c \;+\; \omega_t \;+\; u_{act}. \quad (13)$$

Note that it is appropriate for current, rather than lagged speed, to appear in (13). If a high *lcost_seat_mile* affects utilization independently of speed, then that variable's coefficient should remain negative and significant.

[19] The airline dummy coefficients, which are not reported, show that most carriers fly slower than American, the default carrier. Exceptions are Alaska, Allegiant, United and Virgin American, whose speeds are not significantly different than American's, and Continental, which flies faster. The differences are not great, however, with the largest difference relative to American equal to 1 mph.

Column 4 of Table 3 adds speed as a covariate to the utilization regression in column 2. Recall that if a high *cost_seat_mile* affects utilization independently of speed, then the variable's coefficient should remain negative and significant. This prediction is confirmed, with the cost coefficient smaller than in column 2 but still significant at the 1% level. The speed coefficient is, of course, positive, showing that faster flying yields more available seat-miles. The coefficient magnitude shows that a 1% increase in speed raises available seat-miles by 4.6%. Once again, aircraft-type fixed effects are crucial in identifying these effects.

The results in Table 3 thus show that a high cost per seat-mile reduces an aircraft's flying speed and the available seat-miles it generates. These effects are highly intuitive while also conforming to the predictions of the theoretical analysis.[20]

It is worth asking whether other factors that could affect speed and aircraft utilization are being ignored in our framework, possibly biasing our estimates. For example, a high demand for air travel would presumably raise available seat-miles through an airline supply response. Within the model, high demand could raise the fuel price, thus increasing *cost_seat_mile*. While the resulting effect on available seat-miles would be supply-reducing, the increase in airfares due to stronger demand would encourage more supply. Although failure to capture the latter effect could lead to omitted variable bias in the utilization regression, the presence of year fixed effects in our model precludes such an outcome, with these dummy coefficients taking large values in high-demand years. Therefore, in addition to capturing micro-level aircraft-utilization impacts that operate through *cost_seat_mile*, our framework can capture macro-level demand impacts via its flexible structure, thereby circumventing possible omitted variable bias.

While Brueckner and Abreu (2020) show that fuel consumption naturally varies in 1-to-

[20] A reviewer suggested using carrier×year fixed effects as a way of addressing the possible mismeasurement of fuel prices resulting from time-varying, carrier-specific hedging practices. While this approach is imperfect, adopting it leaves the qualitative results in Table 3 unaffected, with each coefficient increasing slightly in absolute value while retaining significance. The regressions below based on (14), which are run at the airline-year level, do not have enough observations for carrier×year fixed effects, but when this approach is used in the drawdown/buildup regressions based on (15), the results are inferior to those using uninteracted carrier fixed effects. Conceptually, allowing the estimated drawdown and buildup patterns to vary within airlines by year appears to reduce the scope for potential impacts of the explanatory variables, possibly making this approach inappropriate a priori.

1 fashion with available ton-miles (closely related to available seat-miles), they provide no evidence on the speed/fuel-consumption linkage. To generate such evidence, we can estimate their regression (column 3 of their Table 3) on our data while including speed as an additional covariate. The existing covariates in their Table 3, whose coefficients are all significant with the expected signs, are log of available ton-miles (positive and naturally close to 1), load factor (positive), stage length (negative) and the fuel price (negative), along with aircraft-type fixed effects. When their regression is run on our sample without the speed variable, the results are close theirs, with discrepancies due to small differences in the samples. But when the log of speed is included, the stage-length coefficient flips from negative to significantly positive while the *lspeed* coefficient is significant with the wrong negative sign. Apparently, inter-connections between speed, stage length, and ton-miles account for this unexpected result, because when available ton-miles is dropped from the regression, the *lspeed* coefficient switches from negative to significantly positive, indicating that a 1% increase in speed raises fuel use by 6%. However, because this regression is by no means an ideal specification, we do not report it.[21] The upshot is that showing the link between speed and fuel consumption, which we know exists based on engineering evidence, is not straightforward in the presence of other related variables.

5.2. Fleet replacement regressions: Changes in age and fuel efficiency

A carrier makes aircraft replacement decisions several years in advance by placing orders with the manufacturer. In normal times, fulfillment of an order takes just a few years, in contrast to the present-day existence of large order backlogs at Boeing and Airbus, which may lead to long waits for delivery of an aircraft. However, airline leasing of aircraft is extensive, and leasing companies can often provide planes on relatively short notice.

Since fleet replacement swaps old planes for new, more fuel-efficient aircraft, it reduces both the average age and average fuel efficiency of a carrier's fleet. Accordingly, our first approach to analyzing fleet replacement focuses on changes in the average age and fuel effi-

[21] This regression replaces the fuel price, whose coefficient inexpicably turns positive with the omission of available ton-miles, by year fixed effects and adds carrier dummies. In addition to the positive speed effect, this specification leads, as before, to a positive and significant load-factor coefficient, although the stage-length coefficient turns positive and loses significance, again showing the difficulty of crafting a successful specification that includes speed.

ciency (*gallons_seat_mile*) of an airline's fleet. Let Δy_{ct} denote the change in the log of either average aircraft age or average fuel efficiency for airline c between years $t-1$ and t, and let $lfuel\_price\_diff_t$ denote the change in the log fuel price between these years. Since fleet replacement decisions are made well in advance, we also include a lag of the fuel-price difference, given by $lfuel\_price\_diff\_lag_t$, equal to the change in the log fuel price between years $t-2$ and $t-1$. Then the regression equation is

$$\Delta y_{ct} \;=\; \kappa \; lfuel\_price\_diff_t \;+\; \omega \; lfuel\_price\_diff\_lag_t \;+\; X_{ct}'\phi \;+\; \mu_c \;+\; r_{ct}, \qquad (14)$$

where μ_c is again the airline fixed effect, r_{ct} is the error term, and X_{ct} is a column vector of additional covariates.

As for the elements of X, one is a time trend variable equal to $t - 1991$, which is used because year fixed effects are collinear with fuel prices. Two additional variables are the December unemployment rate for the given year (capturing the state of the economy) and a merger dummy, set at 1 in the year after completion of a merger (after the merger partner's aircraft counts become zero).[22] This variable captures a possible change in an airline's average fuel efficiency after absorption of the partner's fleet.

Table 4 shows the results of estimating (14). In column 1, the dependent variable is the first difference of the log of an airline's average gallons per seat-mile, denoted *lavg_gallons_seat_mile _diff*. The regression excludes the current fuel price difference, on the expectation that the lagged difference may be more important. As can be seen, however, the coefficient of this variable, though negative as expected, is insignificant. The same conclusion emerges in column 2, where the dependent variable is the first difference of the log of an airline's average aircraft age, denoted *lavg_age_diff*. The effect of the lagged fuel price difference is again insignificant.

Columns 3 and 4 show the effects of adding the current fuel-price difference to both regressions. As can be seen, the coefficients on the lagged difference remain insignificant, but the

[22] The merger dummy equals 1 for American in 2002 and 2015 following the TWA and US Airways mergers, for US Airways in 2007 following the America West merger, for Delta in 2010 following the Northwest merger, for United in 2010 following the Continental merger, and for Southwest in 2012 following the AirTran merger.

coefficients of the current fuel-price difference are negative and significant, although the significance level in column 4 is slightly below conventional levels, at 6%. Therefore, even though we might expect fuel-price movements several years previously to affect fleet replacement decisions, the regressions show that only the current fuel-price difference (the change between years $t-1$ and t) has a significant effect. A possible reason why the current as opposed to lagged fuel-price difference matters is use of leased planes in fleet adjustments, which allows quicker responses to fuel-price changes than the purchase of new planes.

Overall, the results in Table 4 show that a current fuel-price increase hastens fleet replacement by reducing the fuel consumption and average age of an airline's aircraft. Among the control variables, both the merger and time-trend coefficients are insignificant in all the regressions.[23] But the coefficient of the unemployment variable is significantly positive in the age-difference regressions. Thus, fleet replacement as measured by age changes is evidently slowed (with the changes being less negative) in bad economic times, when the unemployment rate is high.

5.3. Fleet replacement regressions: Drawdown and buildup

The second approach to analyzing fleet replacement focuses on drawdown and buildup of aircraft types. Letting q_{act} denote either of the $drawdown_{act}$ or $buildup_{act}$ variables in (9) and (10), the regression specification is

$$q_{act} \;=\; \xi\, fuel\_price_t + \epsilon\, rel\_gallons_{act} + \upsilon\, fuel\_price_t \times rel\_gallons_{act} + X'_{ct}\phi + \mu_c + w_{act}, \tag{15}$$

where $rel\_gallons_{act}$ is defined in (11), the third term is an interaction term between $fuel\_price$ and $rel\_gallons$, and w_{act} is the error term. Even though we might expect past fuel prices to influence drawdown and buildup decisions given the need to plan in advance, we find that use of the current price in these regressions leads to superior results, mirroring the outcomes seen in Table 4. The reason may again be the influenced of leasing, which allows quicker responses to current conditions.

[23] Adding a squared time trend leaves the qualitative fuel-price effect in column 3 unchanged, while the effect in column 4 becomes significant at a level close to 10%.

Table 5 presents the drawdown and buildup regressions based on (15) with the interaction term omitted, using the unlogged fuel price.[24] Recall that the dependent variable equals the proportional drop in the aircraft-type count for types whose count is falling (drawdown) or the proportional increase in the count for types whose count is rising (buildup). Both measures are thus positive. In the drawdown regression of column 1, the main variables are $fuel\_price$ and $rel\_gallons$, equal to the aircraft type's fuel efficiency relative to the fleet average. As can be seen, the fuel-price coefficient is insignificant while the $rel\_gallons$ coefficient is positive and significant at the 1% level. Therefore, an aircraft type is drawn down faster when $rel\_gallons$ is higher, indicating much worse fuel efficiency relative to the fleet average. But the fuel price appears to play no role in the drawdown process, a conclusion that will be further investigated below. The buildup regression in column 2 shows a mirror-image result, with the significantly negative $rel\_gallons$ coefficient indicating that the buildup of an aircraft type is faster when relative gallons is much lower (fuel efficiency is much better) than the fleet average. The fuel-price coefficient is again insignificant.

These findings are natural, but the absence of fuel-price effects is unexpected. This conclusion is overturned in column 3, where the interaction between $rel\_gallons$ and $fuel\_price$ is added to the regression of column 1. The $rel\_gallons$ effect, which from (15) equals $\widehat{\epsilon} + \widehat{\nu} fuel\_price_t$, remains significantly positive (when evaluated at the mean fuel price for observations with nonzero *drawdown*). However, the overall fuel-price effect, which equals $\widehat{\epsilon} + \widehat{\nu}\, rel\_gallons_{act}$, is again insignificant (when evaluated the mean of $rel\_gallons$). But the positive interaction coefficient indicates that the relative-gallons effect *is stronger when the fuel price is higher*. This finding, which shows that the drawdown of aircraft with higher $rel\_gallons$ is faster the higher is the fuel price, conforms to intuition.

The presence of the interaction term in the buildup regression in column 4 eliminates the significance of all the main coefficients. But the signs of the level coefficients match those in column 2, while the negative point estimate of the interaction coefficient tells the same story as before. In other words, a higher fuel price hastens the buildup of an aircraft type with low

[24] Since drawdowns and buildups involve only a handful of aircraft types, it appears inappropriate to include type fixed effects in the regressions of Table 5. In addition, since the dependent variable in these regressions is not logged, the unlogged fuel price is used on the right-hand side.

relative gallons.[25]

It could be argued that, since the dependent variables are censored at zero in the drawdown and buildup regressions, they should be estimated by a Tobit routine. A counterargument is that the constructed nature of the variables, from (9) and (10), makes the censoring artificial and renders such an approach inappropriate. Regardless of the correct view, Tobit estimation of (15) preserves many of the qualitative features of the results in Table 5, although the interaction coefficient in column 3 loses significance while remaining positive.

Therefore, fuel prices appear to affect airline fleet-replacement decisions in ways that make sense. A faster increase in fuel prices leads to a faster drop in average gallons per mile (a faster improvement in fuel efficiency) and a faster decrease in the average age of a carrier's aircraft. A higher level of the fuel price hastens the drawdown of lower-fuel-efficiency aircraft, while appearing to hasten the buildup of higher-fuel-efficiency aircraft, although the latter effect is insignificant.

6. Where do retired aircraft go?

With the drawdown of older aircraft being an important path to higher fuel efficiency for an airline's fleet, it is natural to wonder where the retired aircraft go. Very often, retired planes are used for crew training or to provide inventories of spare parts, with the latter strategy being particularly profitable if the fleet contains a large number of that type and if the type's drawdown occurs over a number of years.[26] Alternately, retired aircraft can be sold to a leasing company or to another other airline for continued service,[27] with the buyer trading off higher fuel costs for lower capital costs.[28]

Evidence on the dispositions of selected retired aircraft is provided in Table 6. Before considering the numbers, note that the table entries were constructed using individual aircraft histories from Planelist.net and Planespotters.net. Although planes often circulated among a

[25] For a related analysis pertaining to the automobile market, see Klier and Linn (2010), who show that sales of new cars depend on both fuel efficiency and fuel cost, being inversely related to the car model's cost per mile of driving (a variable analogous to our cost PSM).

[26] Alternately, the aircraft could be sold to a parts broker for dismantling.

[27] For an empirical analysis of these resale markets, see Gavazza (2011).

[28] This trade off is particularly favorable for a cargo airline, which may only fly its aircraft once or twice a day, limiting fuel costs relative to more-intensive airline use.

number of different secondary carriers following retirement, aircraft dispositions in the table were assigned based on the predominant use. For example, non-OECD use was assigned when the retired aircraft was mainly operated by passenger carriers in (other) OECD countries even if it was flown by non-OECD airlines for part of its remaining life. The life of the aircraft was calculated as the difference between the year of removal from service and the year it was manufactured.[29]

Turning to Table 6, the first row shows that, at its peak, American Airlines had a large fleet of 270 MD-82 aircraft, which represented 47% of the total world fleet (see the second panel). American retired its fleet over a 10-year period, using the bulk of the retired aircraft (81%) for spare parts or training (some were also donated to museums). Two percent of the aircraft were operated by other OECD passenger airlines, while 12% were operated by non-OECD passenger carriers, with 5% operated by cargo or charter airlines. The first row of the lower panel shows that the ages at removal from service depended on the aircraft's disposition, with longer usage by non-OECD passenger airlines and cargo or charter operators.

United and Southwest retired their vintage 737 aircraft over periods of 7 and 10 years, respectively. United's peak count of 98 737-300 aircraft accounted for 9% of the world total, while Southwest's peak count of 60 737-200 aircraft accounted for 6% the world total. The upper panel of the table shows that a smaller share of these aircraft were retained for parts and training than in case of American's MD-82s. As with American's planes, these 737s were removed from service at greater ages when operated by carriers other than OECD passenger airlines.

Delta's DC-9-30 aircraft (82 planes, accounting for 14% of the world fleet) were released into an expanding low-cost-carrier environment in the US, with the bulk ending up at ValuJet and AirTran before being scrapped. Very few (10%) were retained for parts or training. Regardless of disposition, these DC-9s were removed from service at greater ages than any of the other planes shown in the table.

[29] A small number of aircraft did not have information indicating the date of removal from service, and they were dropped from the data.

7. Conclusion

This paper has documented several important channels by which fuel prices affect fuel usage in the airline industry. Since concerns about climate change make airline fuel usage, and thus aircraft emissions, a central public policy issue, the paper's findings are important. Our results show that, when fuel cost per seat-mile (which depends on both the fuel price and aircraft fuel efficiency) is high, an aircraft type tends to be flown at a lower speed and to generate fewer available seat-miles per year. This negative seat-miles effect is partly due to the lower speed, but our results suggest that fuel-inefficient planes are also used less intensively, spending more time on the ground than their more-efficient counterparts.

The paper also documents a connection between fuel prices and the retirement of inefficient aircraft. A trend of rising fuel prices generates upward and downward trends, respectively, in the average fuel efficiency and average age of an airline's fleet. A similar conclusion emerges for individual aircraft types, with high fuel prices raising the rate at which fuel-inefficient types are drawn down and eventually eliminated from the fleet.

Since airlines do not fully internalize the environmental damage from their fuel consumption, government intervention in the form of an environmental fuel tax is appropriate. Brueckner and Abreu (2017) computed the required magnitude of such a tax, assuming \$40 of environmental damage per metric ton of CO_2, and they reached a value of \$0.39 per gallon of jet fuel. The results of this paper indicate the channels by which such a tax could affect airline operations. Because of the resulting rise in the fuel cost per seat-mile, fuel inefficient aircraft would be flown even slower than they are today and would generate fewer available seat-miles. These inefficient planes would be retired faster than they are today, and the acquisition of more efficient aircraft could be hastened. All these effects would put downward pressure on fuel usage by the airline industry, with consequent environmental benefits.

To derive a quantitative as opposed to qualitative impact, consider the effect on speed of introducing a \$0.39 fuel tax, under the assumption that the pre-tax fuel price is unaffected. The first step is to note that this tax represents 39% of the average fuel price of \$1.005 from Table 1. Next, note that since *cost_seat_mile* rises in proportion to the (after-tax) fuel price, the 39% tax-induced price increase raises this cost by 39%. With a 10% increase in *cost_seat_mile* reducing speed (from above) by 5.7 mph, it follows that a 39% increase will reduce speed by 3.9×5.7 mph, or 22.23 mph, a substantial magnitude. However, given that *cost_seat_mile* is a composite of fuel efficiency and the fuel price, using its coefficient to predict the effect of large fuel-price increases (holding fuel efficiency constant) may give implausible results.[30] The 22 mph number is nevertheless suggestive, indicating that tax policy may be able to substantially affect airline operations and thus fuel usage.

[30] Note that if the change in cost PSM were instead due to a 39% increase in gallons PSM, the speed difference (which now compares planes with substantially different fuel efficiencies) might be more plausible. In the same way, while the model's prediction of a 58.5% drop in available seat-miles from a 39% increase in the fuel price (3.9 times Table 3's 15% reduction from a 10% cost increase) seems implausible, it might not be unreasonable if the 39% cost change were due to the same percentage increase in gallons PSM, implying that dramatically less fuel-efficient planes are much less utilized.

Table 1: Summary statistics

VARIABLES	*Obs.*	*Mean*	*Std. Dev.*	*Min*	*Max*
fuel_price	29	1.001661	.2919667	.6239967	1.550163
gallons_seat_mile	2,058	.0176311	.004464	.0096263	.0429169
cost_seat_mile	2,058	.0164822	.004731	.0085528	.0357984
avl_seat_miles	2,058	1.10e+10	1.18e+10	4553440	9.98e+10
speed	2,058	455.111	46.89648	317.9204	539.8943
avg_gallons_seat_mile	373	.0165887	.0030133	.0111741	.0275682
avg_age	373	9.237527	5.956599	0	32.44
American (AA)	2,058	.127794	.3339414	0	1
Alaska (AS)	2,058	.0461613	.2098853	0	1
Jet Blue (B6)	2,058	.0092323	.0956633	0	1
Continental (CO)	2,058	.0932945	.2909153	0	1
Delta (DL)	2,058	.1686103	.3744984	0	1
Frontier (F9)	2,058	.0272109	.162737	0	1
AirTran (FL)	2,058	.0097182	.0981245	0	1
Allegiant (G4)	2,058	.0155491	.1237528	0	1
Hawaiian (HA)	2,058	.0199223	.1397671	0	1
America West (HP)	2,058	.0335277	.1800537	0	1
Spirit (NK)	2,058	.0199223	.1397671	0	1
Northwest (NW)	2,058	.1015549	.3021355	0	1
TWA (TW)	2,058	.0471331	.2119751	0	1
United (UA)	2,058	.1511176	.3582505	0	1
US Airways (US)	2,058	.074344	.2623937	0	1
Virgin America (VX)	2,058	.0102041	.100523	0	1
Southwest (WN)	2,058	.0447036	.2067026	0	1

Table 2: **Aircraft-Type Frequency and Gallons per Seat-Mile**

Aircraft type	*Frequency*	*Gallons PSM*
Airbus A-318	10	0.0182
Airbus A300-600/R/CF/RCF	19	0.0163
Airbus A300B/C/F-100/200	5	0.0184
Airbus A310-300	3	0.0186
Airbus A319	122	0.0158
Airbus A320-100/200	145	0.0136
Airbus A320-200neo	2	0.0111
Airbus A321-200neo	2	0.0107
Airbus A330-200	36	0.0144
Airbus A330-300/333	16	0.0139
Airbus A330-900	1	0.0131
Airbus A350-900	1	0.0128
Boeing 717-200	13	0.0223
Boeing 727-100	3	0.0315
Boeing 727-200/231A	62	0.0252
Boeing 737-100/200	89	0.0229
Boeing 737-300	120	0.0171
Boeing 737-400	43	0.0167
Boeing 737-500	60	0.0194
Boeing 737-700/700LR/Max 7	74	0.0147
Boeing 737-800	82	0.0130
Boeing 737-900	31	0.0122
Boeing 737 Max 8	4	0.0109
Boeing 737 Max 9	1	0.0100
Boeing 747-100	26	0.0183
Boeing 747-200/300	34	0.0199
Boeing 747-400	34	0.0173
Boeing 747SP	4	0.0252
Boeing 757-200	164	0.0143
Boeing 757-300	34	0.0131
Boeing 767-200/ER/EM	98	0.0178
Boeing 767-300/300ER	110	0.0154

Continued on next page

Table 2 continued

Aircraft type	*Frequency*	*Gallons PSM*
Boeing 767-400/ER	38	0.0150
Boeing 777-200ER/200LR/233LR	79	0.0167
Boeing 777-300/300ER/333ER	10	0.0159
Boeing 787-8 Dreamliner	13	0.0146
Boeing 787-9 Dreamliner	3	0.0124
Boeing 787-10 Dreamliner	1	0.0115
Fokker 100	13	0.0259
Lockheed L-1011-1/100/200	17	0.0195
Lockheed L-1011-500 Tristar	10	0.0229
McDonnell Douglas DC-9-10	19	0.0352
McDonnell Douglas DC-9-30	59	0.0261
McDonnell Douglas DC-9-40	27	0.0266
McDonnell Douglas DC-9-50	28	0.0260
McDonnell Douglas DC-9 Super 80/MD81/82/83/88	144	0.0187
McDonnell Douglas DC-9 Super 87	3	0.0254
McDonnell Douglas DC-10-10	36	0.0183
McDonnell Douglas DC-10-30	45	0.0205
McDonnell Douglas DC-10-40	13	0.0209
McDonnell Douglas MD-11	24	0.0192
McDonnell Douglas MD-90	27	0.0158

Table 3: Speed and Available Seat-Miles Regressions

VARIABLES	(1) lspeed	(2) lspeed	(3) lavl_seat_miles	(4) lavl_seat_miles
lcost_seat_mile	-0.159**	–	–	–
	(0.00780)			
lcost_seat_mile_semilag	–	-0.125**	-1.477**	-0.898**
		(0.00841)	(0.244)	(0.255)
lspeed	–	–	–	4.639**
				(0.677)
Constant	5.522**	5.436**	16.28**	-8.941**
	(0.0446)	(0.0345)	(1.000)	(3.810)
Fixed Effects				
Aircraft type	yes	yes	yes	yes
Airline	yes	yes	yes	yes
Year	yes	yes	yes	yes
Observations	2,058	1,874	1,874	1,874
R^2	0.940	0.942	0.578	0.588

Standard errors in parentheses

** p<0.01, * p<0.05

Table 4: Average Gallons per Seat-Mile and Average Age Regressions

	(1)	(2)	(3)	(4)
VARIABLES	lavg_gallons_seat_mile_diff	lavg_age_diff	lavg_gallons_seat_mile_diff	lavg_age_diff
lfuel_price_diff	–	–	-0.0865**	-0.215†
			(0.0331)	(0.113)
lfuel_price_diff_lag	-0.0288	0.0739	-0.0200	0.0950
	(0.0338)	(0.113)	(0.0337)	(0.113)
unemployment	0.131	2.448*	0.0334	2.198*
	(0.291)	(0.983)	(0.291)	(0.987)
merger	0.000733	-0.124	-0.00108	-0.129
	(0.0322)	(0.106)	(0.0320)	(0.106)
trend	-0.000277	-0.000966	-0.000408	-0.00121
	(0.000716)	(0.00239)	(0.000711)	(0.00238)
Constant	-0.0120	-0.0913	-0.00363	-0.0712
	(0.0256)	(0.0851)	(0.0256)	(0.0854)
Fixed Effects				
Airline	yes	yes	yes	yes
Observations	346	332	346	332
R^2	0.018	0.127	0.038	0.137

Standard errors in parentheses

** p<0.01, * p<0.05, † p<0.06

Table 5: Drawdown and Buildup Regressions

VARIABLES	(1) *drawdown*	(2) *buildup*	(3) *drawdown*	(4) *buildup*
fuel_price	-0.00773	0.0315	-0.172**	0.223
	(0.0147)	(0.0518)	(0.0581)	(0.205)
rel_gallons	0.189**	-0.323**	0.0446	-0.155
	(0.0135)	(0.0476)	(0.0511)	(0.180)
rel_gallons × *fuel_price*			0.162**	-0.189
			(0.0555)	(0.196)
unemployment	-0.328	-0.802	-0.331	-0.799
	(0.201)	(0.707)	(0.200)	(0.707)
trend	0.000803	-0.00297	0.000825	-0.00299
	(0.000487)	(0.00172)	(0.000486)	(0.00172)
constant	-0.126**	0.464**	0.0203	0.294
	(0.0191)	(0.0674)	(0.0536)	(0.189)
Fixed Effects				
Airline	yes	yes	yes	yes
Observations	2,058	2,058	2,058	2,058
R^2	0.111	0.052	0.115	0.052

Standard errors in parentheses

** p<0.01, * p<0.05

Table 6: Fate of Retired Aircraft

Airline	*Aircraft type*	*Peak count*	*Disposition upon leaving fleet*			
			Parts/training	OECD pax	non-OECD pax	Cargo/charter
American	MD-82	270	81%	2%	12%	5%
United	B737-300	98	56%	5%	39%	0%
Southwest	B737-200	60	30%	22%	37%	12%
Delta	DC-9-30	82	10%	54%	26%	11%
		% World	*Average life (years)*			
American	MD-82	47%	24.2	25.6	28.5	29.0
United	B737-300	9%	21.0	21.2	24.6	n/a
Southwest	B737-200	6%	20.8	23.7	28.2	32.6
Delta	DC-9-30	14%	27.9	32.7	29.7	40.2

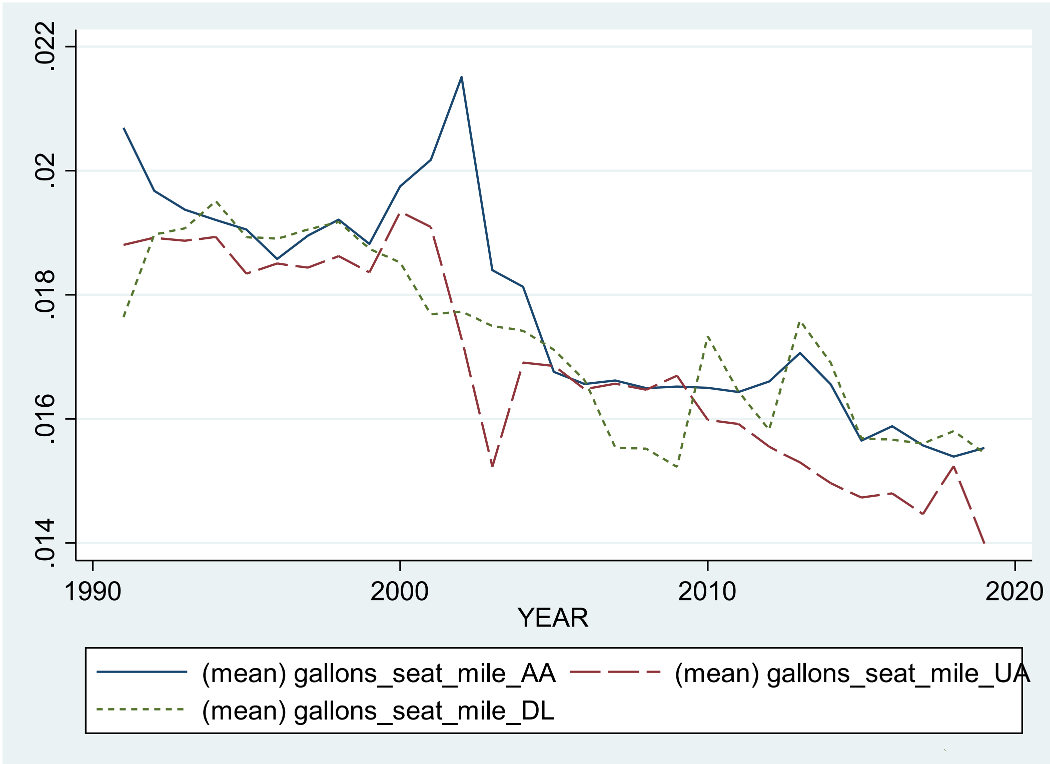

Figure 1: Average Gallons per Seat-Mile by Year for American, United, and Delta

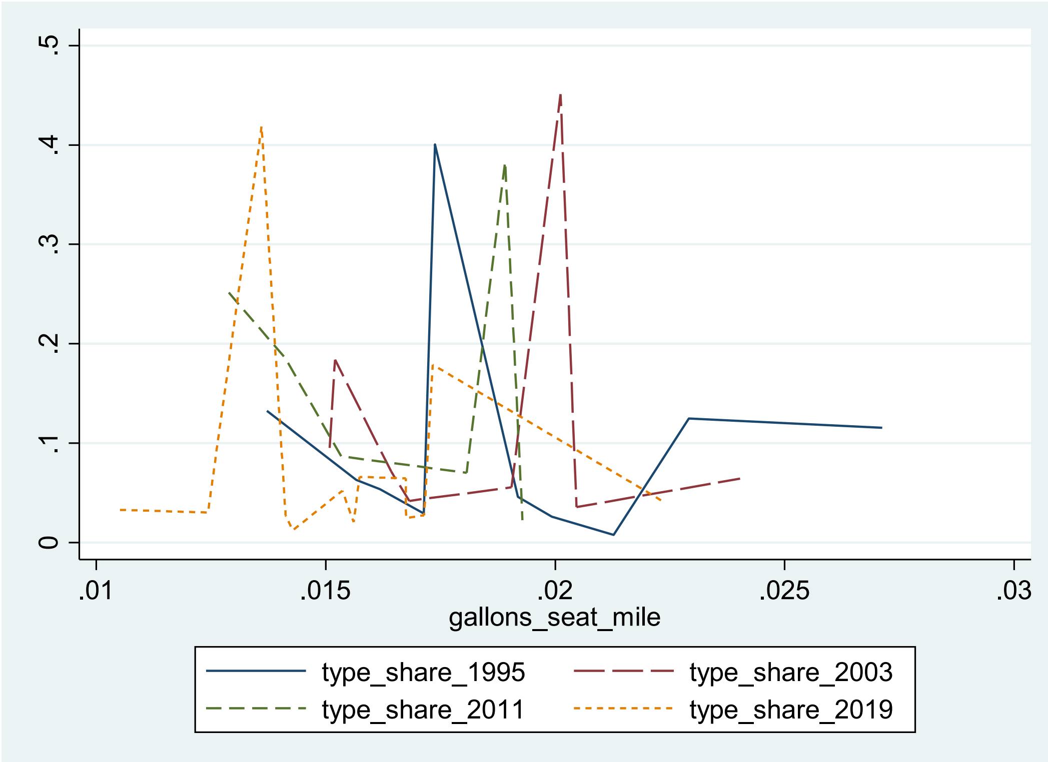

Figure 2: Distributions of American Airlines Aircraft Fuel Efficiency by Year

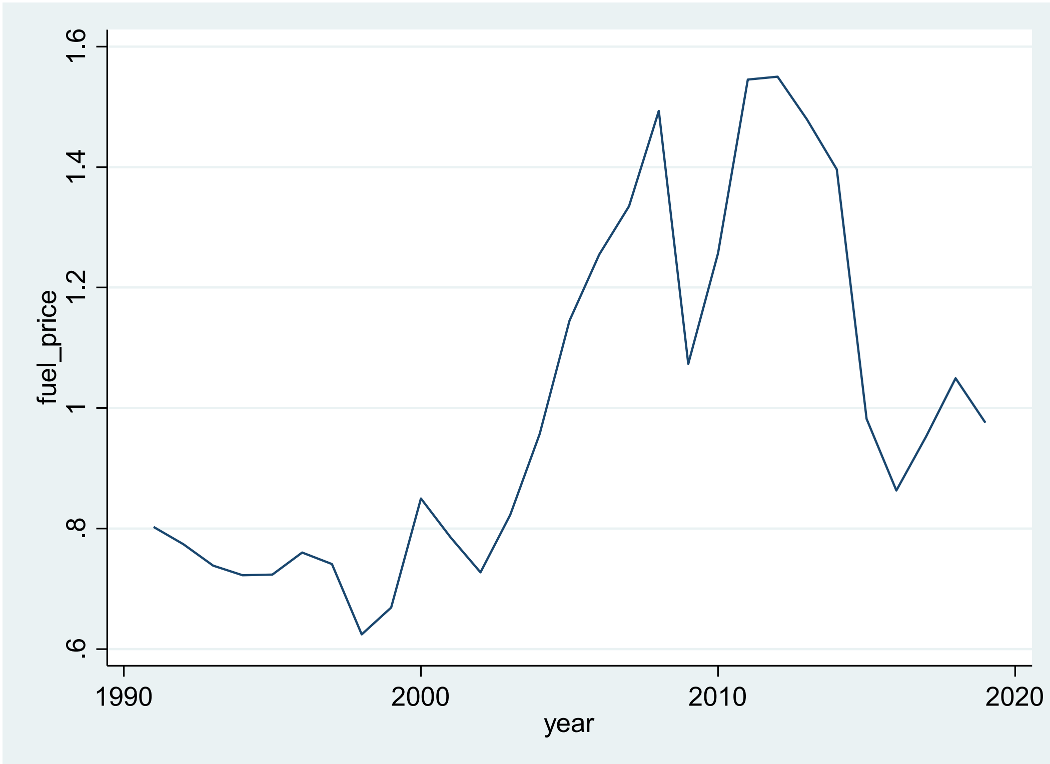

Figure 3: Real Aviation Fuel Price by Year

References

Aktürk, M.S., Atamtürk, A., Gürel, S., 2014 Aircraft rescheduling with cruise speed control. *Operations Research* 62, 829-845.

Boeing, 2007. Fuel conservation strategies: Cost index explained. *AERO Quarterly* 2, 26-28.

Brueckner, J.K., Abreu, C., 2017. Airline fuel usage and carbon emissions: Determining factors. *Journal of Air Transport Management* 62, 10-17.

Brueckner, J.K., Abreu, C., 2020. Does the fuel-conservation effect of higher fuel prices appear at both the aircraft-model and aggregate airline levels? *Economics Letters* 197, article 109647.

de Almeida, E.E., Oliveira, A.V.M., 2023. An econometric analysis for the determinants of flight speed in the air transport of passengers. *Scientific Reports* 13, article 4573. https://doi.org/10.1038/s41598-023-30703-y

Fukui, H., Miyoshi, C., 2017. The impact of aviation fuel tax on consumption and carbon emissions: The case of the US airline industry. *Transportation Research Part D* 50, 234-253.

Fageda, X., Teixido, J.J., 2022. Pricing carbon in the aviation sector: Evidence from the European emissions trading system. *Journal of Environmental Economics and Management* 111, article 102591.

Gavazza, A., 2011. Leasing and secondary markets: Theory and evidence from commercial aircraft. *Journal of Political Economy* 119, 325-377

Goolsbee A., 1998. The business cycle, financial performance, and the retirement of capital goods. *Review of Economic Dynamics* 30, 474-96.

Holland, S.P., Mansur, E.T, Yates, A.J., 2021. The electric vehicle transition and the economics of banning gasoline vehicles. *American Economic Journal: Economic Policy* 13, 316-344.

Kahn, M.E., Nickelsburg, J., 2016. An economic analysis of U.S. airline fuel economy dynamics from 1991 to 2015. National Bureau of Economic Research working paper #22830.

Knittel, C.R., 2012. Reducing petroleum consumption from transportation. *Journal of Economic Perspectives* 26, 93-118.

Li, S., Kahn, M.E., Nickelsburg, J., 2015. Public transit bus procurement: The role of energy prices, regulation and federal subsidies. *Journal of Urban Economics* 87, 57-71.

KLIER, T., LINN, J., 2010. The price of gasoline and new vehicle fuel economy: Evidence from monthly sales data. *American Economic Journal: Economic Policy* 2, 134-153.

MATSUNO, Y., ANDREEVA-MORI, A., 2020. Analysis of achievable airborne delay and compliance rate by speed control: A Case study of international arrivals at Tokyo International Airport. *IEEE Access*. DOI:10.1109/ACCESS.2020.2994109

MERKERT, R., SWIDAN H., 2019. Flying with(out) a safety net: Financial hedging in the airline industry *Transportation Research Part E* 127, 206-219.

MOSKWA, W., 2008. SAS flies slower to save costs and emissions. Reuters, May 20. `https://www.reuters.com/article/us-sas-airspeed/sas-flies-slower-to-save-costs-and-emissions-idUSL2076257020080520`