No 2026-04 – March

Working Paper

CEPII

Air Service Agreements, Connectivity and Emissions

Lionel Fontagné, Cristina Mitaritonna, Gianluca Orefice & Gianluca Santoni

Highlights

- The signature of Air Service Agreements (ASAs) allows airlines to reshape the international route network in a more efficient way and ultimately reduce CO2 emissions per passenger.
- Using unique data on airline tickets and ASAs in force during the period 2012-2019, we show that the re-organization of international flight routes induced by ASAs reduces CO2 emissions per worker by 3.9%.
- The counterfactual analysis suggests that a further liberalization of the air services, in which all country pairs would be linked full liberalization ASAs, would imply a 2.3% reduction in emissions per passenger.

RESEARCH AND EXPERTISE ON THE WORLD ECONOMY

Abstract

The average energy efficiency of the aviation sector has increased by 2.7 percent per year since 2012, falling short of the 6 percent increase in demand. Optimizing routes by reducing the number of legs per flight is one way to complement technological advances in aircraft and fuels to reduce aviation's environmental footprint. The signature of Air Service Agreements (ASAs) allows airlines to reorganize their flight routes. They reshape the international route network in a more efficient way and ultimately reduce CO2 emissions per passenger. On the other hand, ASAs increase the demand for international flights, which may offset the reduction in overall CO2 emissions by airlines. Using unique data on airline tickets and ASAs in force during the period 2012-2019, we show that the considerable reduction in per-passenger CO2 emissions due to the re-organization of international flight routes induced by ASAs is overcompensated by the additional demand for less time-consuming and, hence, more comfortable international flights.

Keywords

Air Service Agreements, Air Transportation, Environment.

JEL

F13, L93, F64.

Working Paper

CEPII

CEPII Working Paper
Contributing to research in international economics

CEPII (Centre d'Études Prospectives et d'Informations Internationales) is a French institute dedicated to producing independent, policy-oriented economic research helpful to understand the international economic environment and challenges in the areas of trade policy, competitiveness, macroeconomics, international finance and growth.

Editorial Director:
Antoine Bouët

Visual design and production:
Laure Boivin

ISSN 2970-491X

March 2026

To subscribe to
The CEPII Newsletter:
www.cepii.fr/KeepInformed

Centre d'études prospectives
et d'informations internationales
20, avenue de Ségur
TSA 10726
75334 Paris Cedex 07

contact@cepii.fr
www.cepii.fr – @CEPII_Paris
Press contact: presse@cepii.fr

1. Introduction

The provision of air services has been shown to affect the economic landscape, facilitate the development of interconnected geographic areas, and encourage face-to-face relationships that support business relationships, international trade, and even technology transfer (Bahar et al., 2023; Blonigen and Cristea, 2015; Campante and Yanagizawa-Drott, 2018; Cristea, 2011; Hovhannisyan and Keller, 2015; Söderlund, 2023). Hence, reducing frictions that hinder the supply of international air services can be beneficial for countries. At the same time, similar to the case of international trade, the environmental impact of air traffic can be divided into technical, scale, and composition effects (Copeland and Taylor, 1993; Grossman and Krueger, 1993; Shapiro, 2025). Air traffic liberalization *via* Air Service Agreements (ASAs hereafter) is expected to have: (i) a *technical effect* (more efficient routes, fewer take-offs and landings as opposed to connecting flights), (ii) a *scale effect* (increased demand for shorter flights) and (iii) a *composition effect* (increased demand for long-distance international flights versus domestic or short-distance flights, and redistribution of market shares among airlines). Understanding the environmental impact of air service liberalization (ASAs) therefore represents a policy-relevant question that remains underexplored in the academic literature.

In this paper, we tackle this question and show how ASAs, by promoting direct flights, shape the carbon emissions footprint of air travel. To do so, we combine granular data on passenger tickets from 2012 to 2019, flight routes, the number of legs, and airport locations with exhaustive information on Air Service Agreements and a detailed simulator of aircraft emissions. Our econometric approach relies on both difference-in-differences techniques and the event study approach to test the parallel trend assumption. By controlling for any country-year and country-pair factors affecting bilateral air services, as well as for any airline-specific technological and productivity shocks, we can claim *causality* in the impact of ASAs on international flight distance, legs, and, hence, emissions per passenger. The reduced-form econometric analysis also informs about the effect of different types of ASAs (depending on their detailed content), showing that full-liberalization agreements are more effective in reducing emissions per passenger. Our policy implications are based on a counterfactual analysis that relies on a structural gravity approach. It accounts for indirect effects driven by the complex interactions between the efficiency of airlines and customer demand for air transport services, and quantifies the overall effect of ASAs on emissions per passenger.

The reduced-form econometric analysis points to some clear-cut results: signing an ASA reduces the distance between origin and destination by 1.6 percent, the number of legs by 3 percent, and the emissions per passenger by 2.8 percent. However, the content of ASAs matters: signing a fully liberalizing ASA reduces emissions per passenger by 5-6 percent, while signing a not-fully liberalizing ASA has no statistically significant impact on emissions per passenger. The counterfactual analysis suggests that the impact of the existing network of ASAs is a 1.9 percent reduction in emissions per passenger, resulting from a 2.8 percent increase in the number of passenger and a 0.9 percent increase in total CO_2 emissions. However, a further liberalization of air services, in which all country pairs would be linked by ambitious (full liberalization) ASAs, would deliver another 2.3 percent reduction in emissions per passenger, a 4 percent increase in demand for more comfortable flights,[1] and a net 1.7 percent

[1]Demand is influenced by the quality (comfort and duration) of flights, but also by ticket prices. Direct flights are generally more expensive per kilometer, as airlines take into account passengers' willingness to

increase in emissions. The main message of the counterfactual analysis is that the scale effect systematically overcomes the technical effect. On the other hand, passenger welfare increases.

This paper builds on three strands of literature. First, discrimination in access to international air services has been shown to be a major obstacle to passenger traffic (Piermartini and Rousová, 2013). The lack of ambition among several ASAs and their network complexity reduce the economic performance of air companies, hinder competition, and eventually increase transportation costs. Open skies agreements and consolidated agreements inspired by those signed by the EU would induce additional passenger traffic. While Piermartini and Rousová (2013) rely on a gravity equation, cross-sectional data on passenger traffic provided by the International Aviation Transport Association (IATA), and an index of the depth of ASAs to assess the impact of sky liberalization on traffic, we contribute by shifting the focus to efficiency in international flight network re-organization and emissions, two key policy-relevant questions. This allows us to assess the scale and efficiency effects of liberalization of passenger transport, and finally, the environmental impact of ASAs. The panel nature of our data (i.e. within identification), combined with specific information on the airlines providing air services, allows us to strengthen the causal interpretation of our findings with respect to previous papers.

A second strand of the literature relies on detailed information on air tickets, flights, or card payments to assess trends in international mobility, traveler consumption behavior, daily emissions, or international trade at the product level. Monthly dyadic air traffic data help to explain the observed trends in increasing international mobility (Gabrielli et al., 2019). The frequency of direct flights and the expansion of airports cause a rise in Chinese travelers' spending abroad, as revealed by card payments (Ho, Chun-Yu and Peng, Tingting and Takayama, Haruka and Xu, Li, 2024). The frequency of passenger flights helps to estimate the impact of in-person interactions on trade at the product level (Wang et al., 2021). Closer to our study, granular information on tickets provides insights into the effect of participating in CORSIA on daily aviation emissions in developing Asian countries (Yan et al., 2025), and on the consequences of the introduction of the EU ETS for aviation on emissions from intra-European flights (Fageda and Teixidó, 2022). We contribute to this specific field of the literature by combining granular information on routes and tickets to estimate the demand and efficiency effects of air service traffic liberalization. The domestic leg of connecting flights, which contributes disproportionately to airline emissions in addition to international flights, is taken into account.[3]

A third strand of the literature assesses the impact of discontinuity in the provision of air transport services. Campante and Yanagizawa-Drott (2018) and Morales-Arilla and Bustos (2024) show that reducing the discontinuities in air transport by opening direct routes for freight would allow countries to specialize and trade, and that this discontinuity impacts the long-term allocation of economic activity in space. Regulatory obstacles that add to technological constraints explain this discontinuity, estimated at 6,000 miles. We contribute by combining geographic information on airports with comprehensive data on the number of flight legs required to cover each origin–destination pair. This represents a substantial contribution,

pay for greater comfort. The presence of an ASA may, therefore, ultimately result in unchanged prices (the shorter distance or flight duration is compensated by a higher price per kilometer), as shown in Table A4 in the Appendix.

[3] Taking domestic flights into consideration is also important for the proper modeling of our counterfactual analysis.

as it allows us to explicitly account for the complex network of direct and indirect flights. Moreover, we focus specifically on the additional demand for passenger transport services induced by more convenient connections and provide a counterfactual analysis showing how the content of ASAs affects the scale and efficiency effects of deregulation.

The remainder of the paper is organized as follows. Section 2 presents the data, provides descriptive statistics, and discusses how we computed the emissions per passenger based on the available data. Our identification strategy is detailed in Section 3. Section 4 combines these elements and provides a reduced-form estimation of the direct environmental impact of ASAs, before quantifying the net effect of ASAS using a structural gravity approach akin to a full-endowment general equilibrium trade model. The last section concludes.

2. Data and Descriptive statistics

To study the impact of air service agreements (ASAs) on airline route organization, we require detailed information on the content of these agreements, ticket purchases (including class) for each origin–destination–year combination, and data on connecting flights, specifically the number of legs needed to reach a destination from a given origin. To this end, we combine three detailed databases: (i) SABRE, which provides ticket information; (ii) Open Flights Airports, which gives the geolocation of terminals used for flights; and (iii) the World Air Service Agreements, which contains the content and signature dates of ASAs. These sources are described in detail in Section 2.1. A final key piece of information concerns aircraft emissions on each route, both with and without connections, which we estimate using a simulator developed at ETH Zurich. By integrating these various elements, we calculate average emissions per passenger for each origin–destination pair on each date – see the detailed discussion in Section 2.2. Finally, some descriptive evidence is presented in section 2.3.

2.1. Primary sources of information

Our primary data source is the air passenger database provided by SABRE, a company that operates a computerized reservation system widely used by travel agencies and airlines to manage flight bookings. SABRE also collects detailed booking and ticketing information from airlines, which is used by these companies to inform their pricing strategies. In this study, we utilize this database for academic purposes.[5] The version of the database used in this study contains annual information on air passenger traffic between 189 countries for the period 2012–2019.[6] This database is unique in terms of both coverage and data quality. It provides air transport information in terms of revenue and passenger numbers, which allows us to infer the average ticket price for flights between two airports. What makes this dataset particularly valuable for our study is that it records the airline selling each ticket, the type of flight (direct or connecting), and detailed information on all intermediate connecting flights from the origin to the final destination, including the airports used for connections. Additionally, the

[5]This data was accessed thanks to a collaboration with, and support from, the Economics of Climate Change, Energy and Transport Unit of the Directorate C - Energy, Transport and Climate of the Joint Research Centre (JRC) in the framework of a report on "The estimation of tariff equivalents for the air service sector" (JRC/SVQ/2020/LVP/1035). The JRC had access to the commercial version of the database and carried out an initial aggregation of the granular data.

[6]We had access to a database aggregated by origin and destination airport, airline, ticket class, and year. In the case of multi-leg connecting flights, we had this information for each leg separately.

database specifies the ticket class (Economy, Economic Premium, Business, or First Class). This information is crucial for calculating passenger CO_2 emissions (see subsection 2.2).

For the baseline analysis, we aggregate the dataset by year, country of origin, and country of destination, while keeping track of the number of flights (legs) required to move from the origin to the final destination.[8] This allows us to compute the year-specific *average* distance needed to travel between any two countries. To construct this measure, we use the Open Flights Airports Database, which provides latitude and longitude coordinates for more than 10,000 airport terminals worldwide.[9] The database also includes ISO country codes and three-letter IATA airport codes, which we use to merge it with the SABRE dataset. Additionally, we employ disaggregated, airline-specific data and estimates to control for airline-specific technological shocks (i.e., fixed effects) when assessing the consequences of ASAs. This approach represents an important contribution to the existing literature.

Our distance measure is primarily based on the geodesic distance between any two airports, accounting for the eventual connections needed. Hence, this variable of distance varies depending on the chosen route, i.e., different hubs for transit and the number of legs. By summing the great-circle distances from a given origin to a given destination airport across airlines and tickets, and dividing by the number of tickets, we obtain the average distance needed to reach a given destination from a given origin. Importantly, since airline companies continuously re-optimize their connecting flight schedules, the average distance varies over time. Similarly, we calculate the average number of flight legs needed to fly from a given origin to a given destination.[10]

Finally, we obtain data on the Air Service Agreements (ASAs) from the World Air Service Agreements (WASA) database, maintained by the International Civil Aviation Organization (ICAO). This is the most comprehensive global database of air service agreements, containing details of signatory countries, the date of signature, and the date the agreement came into force. However, this database comes with incomplete information on the regulation of EU27 countries. We fill this gap by incorporating into the WASA database the set of regulatory reforms experienced by EU27 countries in recent decades.[11] This includes the full liberalization of the EU's internal aviation market since 1992, and the evolution of the EU's external aviation policy since 2002. Indeed, following the indication of the Court of Justice of the European Union (CJEU), any Air Service Agreement (ASA) between an EU Member State and a third country must grant equal market access for routes to destinations outside the EU to any EU carrier with an establishment on its territory.[12] As a result, ASAs between EU Member States and third countries have been amended to incorporate this legal requirement.[13]

[8]We also consider domestic flights, assigning ASA a value equal to zero, as domestic flights do not rely on air service liberalization.

[9]The dataset is available online at https://openflights.org/data.php#country.

[10]SABRE provides data on the total number of passengers for each pair of origin and destination by ticket class (economy, economy premium, business, and first class). Thus, the average distance and number of legs are weighted by the number of passengers flying in each class.

[11]We are indebted to Ana Norman (JRC) for her help and guidance in completing this database.

[12]This has been implemented by countries in two ways. First, through individual bilateral negotiations amending or replacing each ASA separately (Bilateral Air Service Agreements, BASAs). Second, through the European Commission acting on behalf of the EU Member States concerned, negotiating amendments or replacements for multiple preexisting bilateral ASAs with a non-EU country under a single framework (Horizontal Agreements, HAs).

[13]Relevant information from official EU sources was added to the original WASA data to ensure complete

These regulatory changes in EU air transportation have resulted in the amendment of nearly 914 bilateral ASAs between EU Member States and third countries. Of these, 388 agreements were modified through Bilateral Air Service Agreements to recognize the "European nature" of all EU carriers, without changing the original structure of the agreements, while 511 bilateral ASAs were renegotiated as Horizontal Agreements with 36 countries. We merge the amended WASA dataset with the other two primary data sources using ISO country codes.

Importantly, thanks to the information provided by the WASA database, we can classify Air Service Agreements (ASAs) into two categories: (i) *Fully Liberalized Agreements*, including Open Skies agreements, and (ii) *Not Fully Liberalized Agreements*. The Fully Liberalized Agreements contain at least four key liberalization elements: unrestricted traffic rights, multiple designation of airlines (with no route limitations), free determination of capacity (meaning that airlines decide passenger capacity and flight frequency without regulatory restrictions) and dual disapproval or country-of-origin pricing regimes (allowing airlines to set fares freely with minimal government intervention).[15] Not Fully Liberalized Agreements do not include all the liberalization elements listed above.

2.2. Measuring CO2 emissions for aviation passengers

A crucial variable of interest for our study is the level of CO_2 *emissions* allocated to aviation passengers: in total, emissions per passenger (*EP*) and emissions per passenger per kilometer (*EPKM*). Ideally, to accurately estimate these variables, we would need the precise characteristics of the aircraft on which passengers fly. Since this information is not available in our data, we rely on the established methodology employed by online emission simulators. Specifically, we follow the approach developed by the MyClimate Foundation (Myflight, 2024a).

The choice to adopt MyClimate's calculator methodology is based on three key reasons. First, it has a strong scientific foundation. MyClimate was founded in 2002 as a spin-off of ETH Zurich, and its model and parameters are based on estimates from advanced literature and solid statistical analyzes in the field. Second, MyClimate provides comprehensive documentation of the methodology, ensuring transparency and allowing users to understand the assumptions and data underlying the emission estimates. Third, this methodology enables us to calculate the emissions per passenger with greater precision than other available simulators, thanks to the detailed level of flight data in our database. For example, compared to the ICAO Carbon Emissions Calculator (ICAO Carbon Emissions Calculator, 2024) used by Yan et al. (2025), MyClimate allows the calculation of emissions considering both direct flights and flights with stopovers when traveling from an origin to a destination.[16] Furthermore, the parameters of the algorithm used to calculate emissions vary depending on the distance thresholds traveled. Those are crucial aspects of our research.

coverage. All the information needed to amend the WASA database has been recorded through the publications in the Official Journal (OJ) of the European Union and validated by DG MOVE.

[15] Under the "dual" or "double" disapproval method, prices become effective unless both aeronautical authorities disapprove of them. As with the country-of-origin approach (where each party may approve or disapprove of pricing only for flights originating in its own territory) and the free pricing method, the objective is to limit governmental involvement and increase tariff flexibility for designated airlines.

[16] An alternative is the Small Emitters Tool developed by Eurocontrol, which is dedicated to the air traffic covered by the EU emissions trading scheme, i.e. intra-European flights.

So, following MyClimate's methodology, we calculate the CO_2 equivalent emissions per passenger in Kg, by cabin class, for a given flight distance between each origin-destination airport pair by applying the following formula:

$$EP_{class}^{haul} = \left[\frac{a^{haul}x^2 + b^{haul}x + c^{haul}}{S^{haul} \times PLF^{haul}}\right] \times \left[1 - CF\right] \times CW_{class} \times EF \quad (1)$$

where *S* is the number of seats in the flight of a given haul (considering all passenger categories), *PLF* is the passenger load factor,[18] *CF* is the cargo factor,[19] *CW* is the cabin class weighting factor, accounting for the differences in space utilization and fuel consumption per passenger based on seating class,[20] *EF* is the CO_2 emission factor for the combustion of jet fuel (kerosene), which is 3.16 kg CO_2e per kilogram of kerosene.[21] Importantly, parameters *S* and *PLF* depend on the *haul* of the flight: (i) below 1500 Km; (ii) above 2500 Km; and (iii) in between.

The second-order polynomial, $ax^2 + bx + c$, fits the generalized fuel consumption function $f(x) + \mathrm{LTO}$ for any flight of a given haul with distance x, where LTO denotes the extra fuel used during each landing and take-off cycle as well as during the taxi phase (Seymour et al., 2020). Since these factors differ considerably between aircraft types, MyClimate gathers the necessary figures for each variable from the 2019 global data on the most common aircraft. The representative (or standard) aircraft parameters are derived by computing a weighted mean across different models,[22] with separate calculations for short-distance (x $<$ 1500 km) and long-distance (x $>$ 2500 km) routes. For trips spanning 1500 to 2500 km, the values are determined using linear interpolation. See Appendix Table A1 for more details on the parameters used to calculate emissions per passenger.

All these variables are finally merged with our main database to obtain passenger emissions for each pair of origin and destination countries. We first apply the formula in eq. (1) to passengers by class, from the origin airport to the destination airport. For indirect flights with connections, the formula must be applied to each leg of the journey. To obtain the total emissions of passengers flying between an origin country and a destination country, we sum

[18] The methodology considers passenger load factors, recognizing that emissions per passenger decrease as the number of passengers increases, thereby distributing the total emissions among more individuals.

[19] The methodology includes factors such as cargo weight and volume to estimate emissions accurately. Passenger aircraft often transport significant amounts of cargo, particularly wide-body models on long-haul flights. Consequently, a portion of the overall emissions must be assigned to freight. According to the European standard DIN EN 16258 (2012), air freight is now quantified by weight. On international routes, where payloads are larger, cargo emissions are comparatively higher, which in turn lowers emissions attributed exclusively to passengers. The freight factor (CF) is defined as the ratio of cargo weight to total payload. Cargo weight is computed based on the available cargo volume and an average load of 167 kg/m^3. Information on the average available cargo volume, cargo weight, and maximum payload is derived from individual aircraft data sheets.

[20] Typically, the weighting factor is 1 for Economy Class, while it increases for other types of passenger classes due to their greater seat pitch, width, and overall space usage. The MyClimate methodology aligns with international standards such as ICAO and DIN EN 16258 (2012) for fair emission attribution among different passenger classes.

[21] This value is widely used in aircraft emission calculations, including methodologies of organizations such as IATA, or relevant research papers in the field (Lee et al., 2021).

[22] Alternatively, Myclimate provides data on the most and least fuel-efficient variants among the two most prevalent aircraft models used on short-haul and long-haul routes.

the emissions of passengers across all flight segments between the two countries and across all classes of passengers. The variable *EP* is finally obtained by dividing the total emissions by the number of passengers flying between an origin country and a destination country. In contrast, for *EPKM*, we also divide emissions by the distance in kilometers between country pairs.

2.3. Descriptive statistics

Before moving to the econometric analysis, we highlight in this section some key stylized facts about CO_2 emissions, air passenger traffic, and the signature of ASAs during the period 2012–2019.

Fact 1 - A strong increase in the passengers. The total number of air passengers increased by 50% in the period 2012-2019, reaching over 4.4 billion passengers in 2019 (Figure 1). This strong increase suggests the relevance of air service for both policymakers (regulatory policies) and customers. Reassuringly, the SABRE data on air passengers used in this paper align perfectly with the aggregate figure from the website airlines.org (Arlines.org, 2024).[24] Furthermore, the growth in the number of passengers is equally distributed between domestic and international flights. Throughout the period, the former represented 55% and the latter 45%. This increase in the total number of passengers is accompanied by a decrease in the proportion of passengers making multiple stopovers between their point of departure and their final destination – from 32% in 2012 to 26.5% in 2019. Such a reduction in the share of passengers on connecting flights is at the core of our econometric analysis aimed at testing the effect of ASAs on the number of legs and the average distance needed to cover a given origin-destination pair.

[24] Traffic and operations data on this website reflect the scheduled activity of passenger and cargo airlines operating worldwide, as recorded by ICAO.

Figure 1 – Number of passengers, 2012-2019

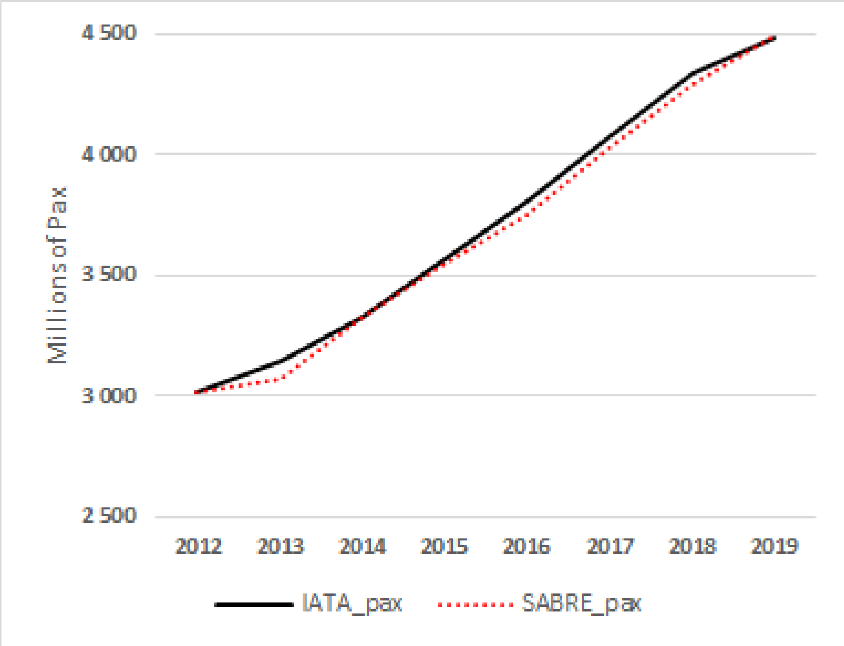

Source: Sabre and ICAO (Arlines.org, 2024)

Fact 2 - The role of ASAs. Between 2012 and 2019, 365 new country pairs signed an ASA. This development reflects a combination of factors: a small number of country pairs signed new non-fully liberalized ASAs (36 country pairs), while 101 country pairs renegotiated existing agreements to make them fully liberalized, leading to a net decline of 65 pairs with not-fully liberalized ASA members. The vast majority of country pairs (228) signed entirely new fully liberalized agreements – see Figure 2a. So, the number of country pairs with fully liberalized ASAs increased by 329 over this period, rising from 16% to 26% of all ASA-covered country pairs by 2019 – see Figure 2a.[26] Interestingly, throughout the entire 2012–2019 period, a constant 80% of international air passengers traveled between countries linked by an ASA. In contrast, passenger traffic on routes between countries with fully liberalized ASAs increased by over 72%, significantly outpacing the international average – see Figure 2b.

Fact 3 - The CO_2 emission of international flights. Global aviation growth has led to an increase in CO_2 emissions associated with the sector. According to IATA, global air passenger aviation emitted over 820 million tons of CO_2 in 2019 (Figure 3a), accounting for 2.2% of total global CO_2 emissions (Figure 3b).[27] Reassuringly, the calculation algorithm used in this paper

[26] The proportion of country pairs with any ASA relative to all possible pairings grew modestly, from 18% in 2012 to 20% in 2019.

[27] More precisely, Edgar (2024) 's statistics display emissions for the civil aviation sector, encompassing both passengers and cargo: 1036 million metric tons of CO_2 in 2019. The reference research paper in the field yields a very similar amount (Lee et al., 2021). The CO_2 emissions attributed exclusively to passengers in Figure 3a are estimated by deducting cargo emissions from the total emissions of the civil aviation sector, assuming a cargo factor of 20%, based on revenues (IATA, 2019).

Figure 2 – The evolution of ASAs agreements and number of passenger by type of ASAs. Period 2012-2019.

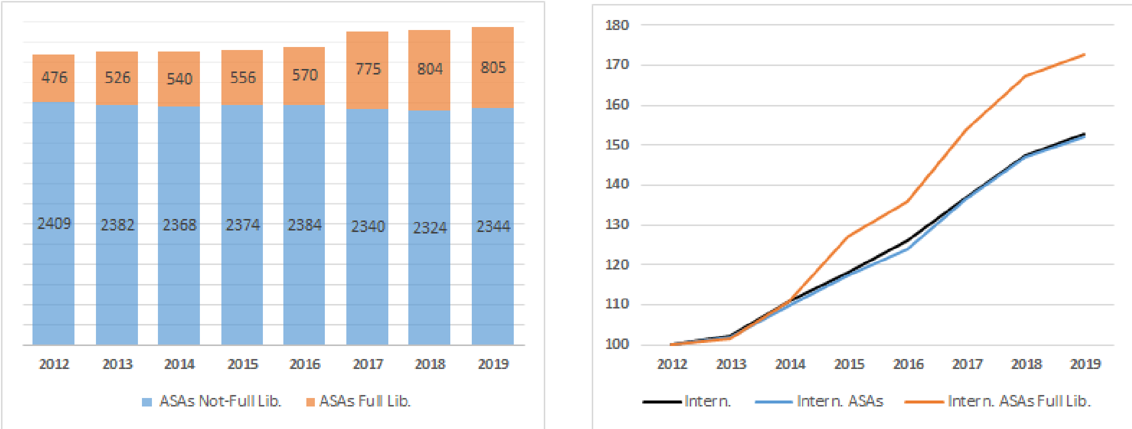

(a) County-pairs sharing ASAs agreements, by type.
Source: Sabre

(b) Evolution of International Air passengers, (2012=100).
Source: Sabre

to infer the emissions related to air passengers (discussed in section 2.2) produces comparable values, i.e., 828 million metric tons CO_2 for 2019.

In the period 2012-2019, the total CO_2 emissions increased at a lower rate (32%) than the total number of passengers (50%), leading to a reduction in the amount of emissions per passenger, EP – see Figure 3a. Such a reduction in the EP in the period 2012-2019 corresponds to an average annual decline of 2.7% between 2012 and 2019,[29] and may be due to several factors: technological advancements in aircraft, better fuel quality, larger average aircraft size, and a concurrent increase in passenger load factors. The primary motivation for this study is to evaluate the extent to which the signature of an ASA may also have contributed to the reduction in EP, particularly by investigating their effect on the optimization of the route and the reduction of flight legs.

Figure 3 – Air passengers CO2 emissions. Period 2012-2019.

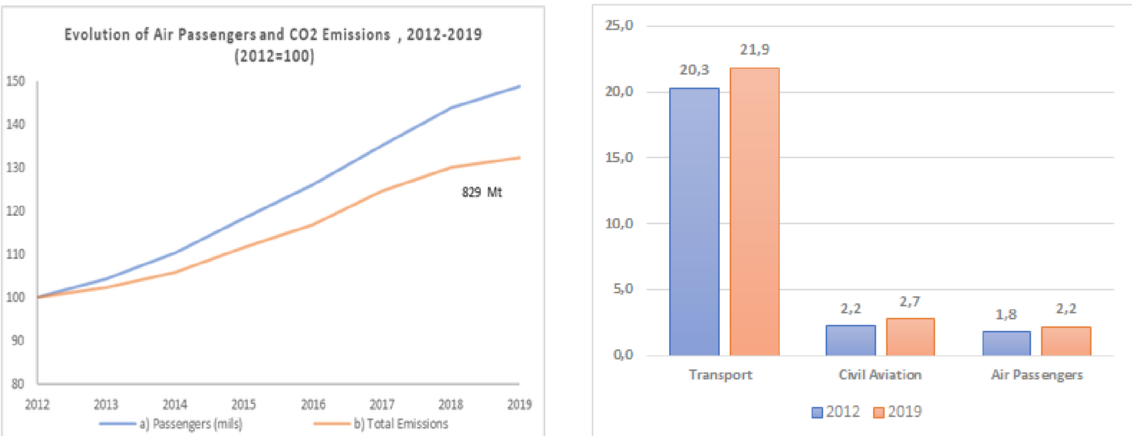

(a) Evolution of Air passengers CO2 emissions, (2012=100)
Source: IATA (Edgar, 2024) and ICAO (Arlines.org, 2024)

(b) Share of CO2 Emissions by type over total.
Source: IATA (Edgar, 2024)

[29]The average annual decline in EP has been computed as $((CO2_{2019}/CO2_{2012})*\exp(1/(n\text{-}1)\text{-}1)*100$.

3. Empirical strategy

We use two complementary reduced-form approaches to quantify how ASAs affect aviation-related outcomes. First, using the nonlinear difference-in-differences (DiD) framework developed in Wooldridge (2023) – which is particularly well suited to data with a gravity-like structure – we estimate the average impact of signing an ASA. In particular, we follow the application of structural gravity models introduced by Nagengast and Yotov (2025) (see section 3.1). We then refine our analysis using an event-study approach to uncover the dynamic effects of ASAs and test the parallel trends assumption (see section 3.2).

3.1. Difference-in-Difference approach

We first estimate the impact of signing an ASA on air traffic and emissions per passenger using the following nonlinear difference-in-differences approach:

$$y_{odt} = exp\left[\theta_{od} + \theta_{dt} + \theta_{ot} + \sum_{g=q}^{T}\sum_{s=g}^{T}\delta_{gs}ASA_{gs}\right] \times \varepsilon_{odt} \quad (2)$$

where the subscripts *o* and *d* indicate, respectively, the origin and destination of air traffic-related outcomes in year *t*. The main outcome of interest y_{odt} is the amount of emissions per passenger (EP) and emissions per passenger-kilometer ($EPKM$) as defined in section 2.2. While emissions per passenger intuitively measure the environmental efficiency of flights, rescaling by distance (in kilometers) between origin and destination captures differences in emissions profiles between long-haul and short-haul flights.

To identify the channels through which emissions per passenger are affected, we also test the effect of ASA agreements on (i) the number of passengers flying from o to d, (ii) the number of stopovers (legs), and (iii) the average distance traveled between o and d. By enabling direct connections between countries, ASAs are expected to reduce stopovers and travel distances. Simultaneously, by shortening travel times and potentially lowering prices, such agreements may increase passenger volumes.

The origin-year fixed effects (θ_{ot}) and destination-year fixed effects (θ_{dt}) control for country-specific, time-varying factors affecting air traffic (e.g., supply and demand shocks, technological advances, and infrastructure improvements). The presence of country-pair fixed effects (θ_{od}) implies identification from within-pair variation in ASA adoption over the period 2012-2019. The indicator ASA_{gs} is equal to one if a country pair od belonging to the treatment cohort g is in a post-treatment period s (i.e., $s \geq$ ASA adoption year), and zero otherwise.[31] This specification enables equation (2) to capture heterogeneous treatment effects that vary both across cohorts and over time, as represented by the parameters δ_{gs}. In this case, a traditional two-way fixed effects (TWFE) estimator would yield biased estimates if treatment effects vary significantly across units and over time.[32] In linear settings, several heterogeneity-robust

[31]A country pair od belongs to treatment cohort g if the signature of the ASA was in year g. q is the first year of the treatment of cohort g, and T is the last year of the panel.

[32]This bias may result from inappropriate weighting of group-time average treatment effects - including negative weights - and the use of previously treated units as controls, both of which arise when TWFE estimators are applied in the presence of treatment effect heterogeneity across cohorts or over time (Borusyak

DiD approaches have been developed (see, for example, the surveys by De Chaisemartin and d'Haultfoeuille (2023) and Roth et al. (2023)). In nonlinear settings, Wooldridge (2023) derives an estimator that allows for cohort-time-specific treatment effects under the assumption of nonlinear parallel trends.[33] As presented in Nagengast and Yotov (2025), the aggregate treatment effect (β^{ASA}) from equation (2) can be expressed as a weighted average of the cohort–time–specific treatment effects:

$$\beta^{ASA} = \sum_g \sum_s \frac{N_{gs}}{N_D} \hat{\delta}_{gs}$$

where N_{gs} denotes the number of treated observations in cohort g and period s, N_D is the total number of treated observations, and $\hat{\delta}_{gs}$ is the estimated treatment effect for cohort g in event period s. In our difference-in-differences estimations, we systematically use two alternative control groups: non-treated, *versus* non-treated and not yet treated. While β^{ASA} is expected to be negative for outcomes such as the number of stopovers and the average travel distance, and positive for the number of passengers, its sign for emissions per passenger depends on efficiency gains and composition effects. Finally, if the efficiency effect persists even when the distance traveled is shorter (fewer takeoffs and landings), then β^{ASA} is expected to be negative for emissions per passenger per kilometer (EPKM).

Since our empirical approach is inspired by the structural gravity models often used in the international trade literature, we also present the aggregate ASA effect (β^{ASA}) within a typical gravity framework (TWFE estimator) as a benchmark. In all the regressions, we employ a Poisson pseudo-maximum likelihood (PPML) estimator to mitigate heteroskedasticity concerns (Santos-Silva and Tenreyro, 2006) and cluster standard errors two-way at the origin and destination levels to address serial correlation in country-specific outcomes (e.g., passenger flows, emissions).[35] Furthermore, we restrict our sample to commercial flights (IATA) and drop outlier observations (i.e. airline-year combinations below the bottom 1 percentile in total number of passengers).[36]

Although the extensive sets of fixed effects reduce concerns regarding omitted variables, endogeneity due to preexisting bilateral trends that affect ASA adoption remains a concern. To address this issue, the following section introduces an event–study design to capture dynamic ASA effects and to validate the parallel trends assumption.

3.2. Event-study approach

To examine the dynamic effects of ASAs and to validate the parallel trends assumption, we implement an event-study design based on equation (2). In line with Sun and Abraham (2021), we restrict our analysis to comparing origin–destination pairs that have entered into an ASA with those that remain untreated (i.e., never treated) throughout the sample period. This

et al., 2024; De Chaisemartin and d'Haultfoeuille, 2020).

[33]This extends the approach developed in Wooldridge (2021) to nonlinear models.

[35]In-sample descriptive statistics of our variables of interest are reported in Appendix Table A2.

[36]This cleaning is needed as raw data contain many airline-year combinations with a suspiciously small number of passengers, often equal to one. We consider these imprecise data and hence drop them. We show robustness checks of our results using the full sample (i.e. including non-commercial flights and the bottom 1 percentile observations) in Appendix section A.

approach avoids potential bias arising from comparisons with units that are treated at later dates. Following Nagengast and Yotov (2025), our event-study estimator is defined as:

$$\beta_s^{ASA} = \sum_g \frac{N_{gs}}{N_{\cdot s}} \hat{\delta}_{gs},$$

where $\hat{\delta}_{gs}$ denotes the estimated treatment effect for cohort g in event period s, N_{gs} is the number of treated observations in that cohort–period, and $N_{\cdot s}$ is the total number of treated observations in period s. This weighted average captures the dynamic treatment response by ensuring that each event period's effect is appropriately scaled by the number of treated pairs. Under the parallel trends assumption – i.e., in the absence of treatment, the evolution of outcomes would have been identical across treated and untreated pairs – the coefficients $\hat{\delta}_{gs}$ capture the causal impact of ASAs at each relative time period. This dynamic specification not only allows us to trace the evolution of the ASA effect over time but also serves as a diagnostic tool by testing whether pre-treatment estimates are statistically indistinguishable from zero.

3.3. Heterogeneous ASAs

As discussed above, Air Services Agreements contain several types of provisions, and the effect of ASAs on air service-related outcomes may depend on their depth (i.e., content). Hence, we replicate the analysis discussed in sections 3.1 and 3.2 by type of ASAs. Namely, we define *Full Liberalization* ASAs as those agreements that contain at least four key liberalization elements: unrestricted traffic rights, multiple designations of airlines with no route limitations, free determination of capacity, and dual disapproval or country-of-origin pricing regimes. Conversely, we define *Not-Full Liberalization* ASAs as those agreements that do not include all these liberalization elements. Differentiating the effects of these two broad categories of ASAs is informative for policymakers involved in the design of future ASAs.

4. Quantifying the environmental impact of ASAs

In this section, we analyze the effect of ASAs on outcomes related to air services, focusing on emissions per passenger, average distance, and the number of legs required to travel from an origin o to a destination d. We first present our empirical reduced-form estimates using difference-in-differences and event study designs –Section 4.1. We evaluate in Section 4.2 the net effects of ASAs using a structural gravity approach.

4.1. Reduced-form estimation of the impact of ASAs

Table 1 reports the results from a heterogeneity-robust DiD specification based on eq. (2). In panel (a) we show results using robust diff-in-diff estimator excluding not yet treated od pairs, in panel (b) we use the same estimator but consider both not yet and never treated pairs as control group, in panel (c) we use a standard TWFE estimator. Column (1) shows the effect of ASAs on total airline emissions conditional on the number of passengers transported.[38] In

[38] In practice, this is a constrained regression in which total emissions are the dependent variable and log passenger numbers are used as an offset. This imposes a unit elasticity of emissions with respect to passengers,

columns (2) and (3), we consider respectively emissions per passenger (EP) and emissions per kilometer (EPKM).[39] Columns (4) and (5) focus on flight (average) distance and the number of stopovers.

All results are statistically significant and robust across alternative estimators, suggesting that ASAs systematically improve the environmental efficiency of international flights. Following the signing of an ASA in each origin-destination pair, emissions per passenger decrease by 3.9% in our preferred specification (column 2), where the control group consists of country pairs that have never been treated and not-yet been treated, and by 2.8% when the control group consists of never-treated country pairs.

While column (2) treats all routes equally because the dependent variable is emissions per passenger without constraint, column (1) shows the results of a constrained regression, imposing a unit elasticity of emissions with respect to passengers. Therefore, in column (1) the impact on total emissions, conditional on the number of passengers (i.e. efficiency gain per passenger), is slightly smaller, at 2.4 and 2.5% respectively: important routes carry more passengers. In contrast, the larger estimate of -3.9% indicates that more marginal routes experience stronger efficiency gains, consistent with the effects of route reorganization and passenger reallocation.

However, emissions per passenger outcomes do not account for the distance traveled or for differences between short- and long-haul flights.[41] To address this limitation, we use EPKM – i.e., we divide in column (3) total emissions by both the number of passengers and the distance traveled between an origin and a destination country. This enables us to evaluate whether the efficiency gains from ASAs are robust when accounting for varying flight distances. Results show that the signature of an ASA reduces emissions per passenger per Kilometer by 5.5% in our preferred specification. Importantly, point estimates reported in columns (2) and (3) of Table 1 are not statistically different, suggesting that using EP does not introduce bias due to differences between short- and long-haul flights. Therefore, we focus on emissions per passenger and total emissions constrained by the number of passengers in the following sections.

Columns (4) and (5) show that such a reduction in emissions per passenger is driven by shorter routes and fewer stopovers: the average distance decreases by 1.4–1.6%, and the number of legs decreases by 3–4%. These results indicate that establishing an agreement induces airlines to reorganize their routes, thereby reducing travel distances and the number of stops, which, in turn, lowers emissions per passenger. We will discuss in the next sub-section the interpretation of the overall efficiency gain suggested by the results in Table 1: a technical effect is at play (fewer takeoffs and landings, different aircraft) potentially combined with a composition effect (more efficient airlines benefiting disproportionately from ASAs, or specific airlines achieving more efficiency gains).

It must be noted that TWFE and heterogeneity-robust DiD estimators deliver similar, not statistically different, results. This suggests that, in our specific case, the effect of ASA treatment is fairly stable across cohorts and over time, and the traditional two-way fixed

holding scale fixed.

[39]The dependent variable is defined as total emissions divided by either the total number of passengers in column (2) or the product of passengers and kilometers in column (3), without imposing any restriction.

[41]Longer flights are typically more fuel-efficient per kilometer, whereas short flights have higher emission intensity due to the fuel consumed during takeoff and landing.

effects estimator does not underestimate the effects of ASAs on emissions per passenger.[42] Hence, we can fairly rely on TWFE estimates in the structural gravity approach reported in Section 4.2 to quantify the net effect of ASAs.[43]

Table 1 – The impact of ASAs on CO2 emissions, average distance and number of legs.

Dep Var. :	Tot. emissions (1)	EP (2)	EPKM (3)	Distance (4)	# legs (5)
Estimator	Panel (a): Heterogeneity-robust DiD (never treated)				
ASA	-0.025***	-0.028***	-0.041***	-0.016***	-0.030***
	(0.010)	(0.0095)	(0.010)	(0.005)	(0.008)
Obs.	138,310	138,310	138,310	138,310	138,310
Estimator	Panel (b): Heterogeneity-robust DiD (never- & not-yet-treated)				
ASA	-0.024***	-0.039***	-0.055***	-0.014***	-0.039***
	(0.006)	(0.011)	(0.012)	(0.004)	(0.009)
Obs.	138,310	138,310	138,310	138,310	138,483
Estimator	Panel (c): Structural gravity (TWFE)				
ASA	-0.014	-0.026*	-0.047***	-0.012**	-0.042***
	(0.010)	(0.0013)	(0.0014)	(0.005)	(0.010)
Obs.	176,638	176,638	176,638	176,638	176,638
Origin-Year FE	Yes	Yes	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes	Yes	Yes
Origin-Destination FE	Yes	Yes	Yes	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variables are defined as follows: column (1) reports total emissions with log passenger numbers included as an offset, which constrains the elasticity of emissions with respect to passengers to one ("passenger-constrained"); column (2) reports emissions per passenger (EP); column (3) reports emissions per passenger-kilometer (EPKM); column (4) reports the average distance traveled between origin and destination; and column (5) reports the average number of flight segments. Standard errors are two-way clustered at the level of origin and destination countries. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. When the control group is restricted to never-treated country pairs, all pre-treatment observations of not-yet-treated pairs are absorbed by cohort-year fixed effects; this procedure leaves the number of observations unchanged.

While the signature of ASA could reduce the administrative costs for airline companies to connect a given origin-destination pair and induce downward pressure on average prices, the reduced distance and number of stopovers make flights more comfortable, which may allow companies to charge higher prices per kilometer. In Appendix Table A4 we test the effect of ASA on average prices and price per kilometer.[45] Interestingly, the two opposing channels mentioned above offset each other (i.e. null effect of ASA on the average price), but the price per kilometer increases in the presence of ASA, suggesting that companies absorb part of the additional consumer surplus associated with more comfortable and shorter flights.

Finally, the content of ASA matters. In Table 2 we show the effect of full-liberalization vs non-full liberalization ASAs. While full-liberalization ASAs reduce emissions per passenger by 6% – see column (2) in Table 2, non-full liberalization ASAs have no statistically significant

[42] Nagengast and Yotov (2025) suggest that a standard TWFE estimator may produce downward biased results if the effect of the treatment varies across cohorts and over time. This does not seem to be the case in our empirical framework.

[43] In Appendix Table A3 we show robustness checks of our baseline results using the full sample (i.e. including non-commercial flights), and our results hold.

[45] Table A4 also reports results on total revenues for the interested reader.

impact on emissions per passenger. These results are robust to using the full sample, including non-commercial flights (see Appendix Table A5). In Appendix Table A6 we also show the effect of heterogeneous ASAs (i.e. full vs non-full) on total emissions and the number of passengers used in the quantification exercise in 4.2.

Table 2 – The impact of ASAs on CO2 emissions per passenger by type of ASA.

Dep Var.	Emissions per passenger (EP)		
	(1)	(2)	(3)
ASA Full lib.	-0.054***	-0.060***	-0.051***
	(0.012)	(0.013)	(0.015)
ASA Not-Full Lib.	0.004	-0.011	0.005
	(0.012)	(0.011)	(0.015)
Obs.	138,310	138,310	176,638
Estimator	Heterogeneity-robust DiD		Structural
	never-treated	never- & not-yet-treated	gravity (TWFE)
Origin-Year FE	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes
Origin-Destination FE	Yes	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variable is emissions per passenger (EP). Standard errors are two-way clustered at the level of origin and destination countries * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. When the control group is restricted to never-treated country pairs, all pre-treatment observations of not-yet-treated pairs are absorbed by cohort-year fixed effects; this procedure leaves the number of observations unchanged.

4.1.1. Controlling for composition effects

While country–year and country-pair–specific unobserved factors are captured by the fixed effects in our baseline estimations, average improvements in aircraft efficiency or other airline-specific technological advancements are not. If such improvements are correlated with the timing of ASA signatures, their omission could bias our baseline results. The same remark applies to a redistribution of traffic towards more efficient airlines on a given route, leading to a composition effect. Indeed, to the extent that ASAs facilitate long-distance flights, airlines may respond by adopting more efficient aircraft technologies tailored to long-haul operations. Consequently, the overall emissions effects of ASAs reported in Tables 1 and 2 capture not only changes in average distance and legs (i.e. re-organization effect) but also such compositional adjustments in fleet technology. In this section, we address this issue by explicitly controlling for the additional competition and technological channels that may account for part of the compositional effect.

The richness of our data allows us to conduct this additional test. Specifically, we observe the airline associated with each passenger ticket, which enables us to include airline–year and origin–destination–airline fixed effects. The inclusion of airline–year fixed effects in Table 3 controls for firm-specific unobserved technological shocks that may be contemporaneous with the signing of an ASA. This represents a key advantage of our data and identification strategy, as it further mitigates concerns about omitted variables. In column (1), we consider total airline emissions while conditioning on the number of passengers transported; in column (2), the dependent variable is emissions per passenger.[47] The results from these highly demanding specifications confirm our earlier findings: ASAs reduce emissions per passenger.

[47] In Table 3 we focus on the never- and not-yet-treated control group.

Column (1) reports the change in average airline emissions following the signing of an ASA, conditional on the number of passengers transported within an origin–destination–airline triplet. Emissions decrease by 1.0%. This result can be compared with the 2.5% reduction of emissions observed in column 1 of Table 1 using the same control group. Based on this comparison, we can conclude that almost half of the overall impact of ASAs on emissions is channeled through a composition effect, and the other half through a technique effect: airlines shifted towards more efficient aircraft on the newly opened routes, while different airlines benefited unevenly from ASAs, implying larger point estimates in aggregated estimations where firm-specific technological shocks are not accounted for. Column (2) shows a 1.4% reduction in emissions per passenger in the presence of an ASA.

Table 3 – The impact of ASAs on CO2 emissions. Airline-level estimations.

Dep Var.	Total emissions (1)	EP (2)
ASA	-0.010**	-0.014**
	(0.004)	(0.007)
Obs.	1,555,722	1,555,722
Estimator	Heterogeneity-robust DiD never- & not-yet-treated	
Origin-Year FE	Yes	Yes
Destination-Year FE	Yes	Yes
Origin-Destination Airline FE	Yes	Yes
Airline-Year FE	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variables are defined as follows: column (1) reports total emissions with log passenger numbers included as an offset, which constrains the elasticity of emissions with respect to passengers to one ("passenger-constrained"); column (2) reports emissions per passenger (EP). Standard errors are three-way clustered at the origin, destination, and airline-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Appendix Table A7 we show a battery of robustness checks of airline-level results. First, in column (1) of Table A7 we show a check using the full sample. In column (2) we use IATA flights only but keep bottom 1 percentile companies in the sample. In column (3) we keep flights with distance greater than 1,500 Km for which the train is not a valuable option. Finally, in column (4), we exclude national flights. Our results are robust to all these checks.

4.1.2. The heterogeneous effect of ASA

Results so far show that Air Service Agreements provide airline companies with the opportunity to reach new destinations with direct flight (or fewer stopovers) and benefit from a more efficient route schedule. However, the composition effect uncovered in the previous section suggests the uneven technical improvements by airlines in the presence of ASAs: not all airline companies have the supply and financial capacity to benefit from the opportunities offered by the presence of an ASA.

Small, low-profit airlines are often capacity-constrained in reaching new destinations opened by ASAs and are financially constrained in obtaining slots at new destination airports and in eventually rescheduling flight routes. On the contrary, large and highly profitable airline companies have the fleet and financial capacity to benefit from the ASAs: they can afford the fixed costs of obtaining slots in new destinations' airports and more easily re-schedule their offerings on more efficient routes. Hence, in this section, we test the heterogeneous effect of ASAs on airlines of different types: (i) large vs small, and (ii) high- vs low-profitable companies.

We start by showing in Table 4 the differential effect of ASAs (of any type) on airline companies belonging to the top-25 percentile in total 2012's revenues (here used as a proxy of companies' size), and alternatively on top-50 companies in the Skytrax ranking (here used as a proxy of companies' profitability).[50] Results in Table 4 show a first element of heterogeneity in the effect of ASAs: only large and high-profitable companies benefit from ASAs and reduce their emissions per passenger in columns (1) and (2). The estimations using the Top-50 companies dummy as a proxy of profitability confirms these findings - see columns (3) and (4). In Table 4 we always control for company-year fixed effects and origin-destination airline fixed effects, so results are conditional to any company-specific technological shock.

Table 4 – The heterogeneous effect of ASAs (any type) on CO2 emissions.

Dep Var.	Total emissions (1)	EP (2)	Total emissions (3)	EP (4)
ASA × Top-25 Rev t_0	-0.032***	-0.031***		
	(0.062)	(0.0011)		
ASA × Non Top-25 Rev t_0	-0.0007	0.0065		
	(0.0035)	(0.0074)		
ASA × Top-50 Skytrax t_0			-0.029**	-0.029***
			(0.006)	(0.012)
ASA × Non Top-50 Skytrax t_0			0.0004	-0.004
			(0.004)	(0.007)
Obs.	1,555,722	1,555,722	1,555,722	1,555,722
Estimator	Heterogeneity-robust DiD never- & not-yet-treated			
Origin-Year FE	Yes	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes	Yes
Origin-Destination Airline FE	Yes	Yes	Yes	Yes
Airline-Year FE	Yes	Yes	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variables are defined as follows: columns (1) and (3) report total emissions with log passenger numbers included as an offset, which constrains the elasticity of emissions with respect to passengers to one ("passenger-constrained"); columns (2) and (4) report emissions per passenger (EP). Standard errors are three-way clustered at the origin, destination, and airline-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.1.3. Dynamic effect of ASA

The dynamic effect of ASAs on the outcomes shown in columns (1)-(5) in Table 1 is reported in Figure 4. Two key results emerge from the event-study approach. First, for all the outcomes of interest, the parallel-trend assumption holds: the patterns for treated and untreated country pairs do not differ in the pre-ASA period.[51] This reassures us on the causal interpretation

[50] https://www.worldairlineawards.com/the-worlds-top-100-airlines-2012/. The ranking of airlines is based on a survey of consumers' favorite companies, reflecting their popularity and brand image, and potentially their profitability.

[51] There is a slightly significant pre-trend in the panel (d) of Figure 4 when the number of legs is used as an outcome variable.

of our baseline results. Second, the effect of ASA on emissions per passenger is negative already in the year of the signature, and becomes stronger one and two (and more) years after the signature – see panels (a)-(b) of Figure 4. The same pattern emerges from the average distance reported in panel (c) of Figure 4. Concerning the number of legs, there is no further reduction, as expected, after the initial adjustment to the new possibilities offered by the signed agreement – see panel (d) of Figure 4.

Figure 4 – The dynamic effect of ASAs. Event-study approach.

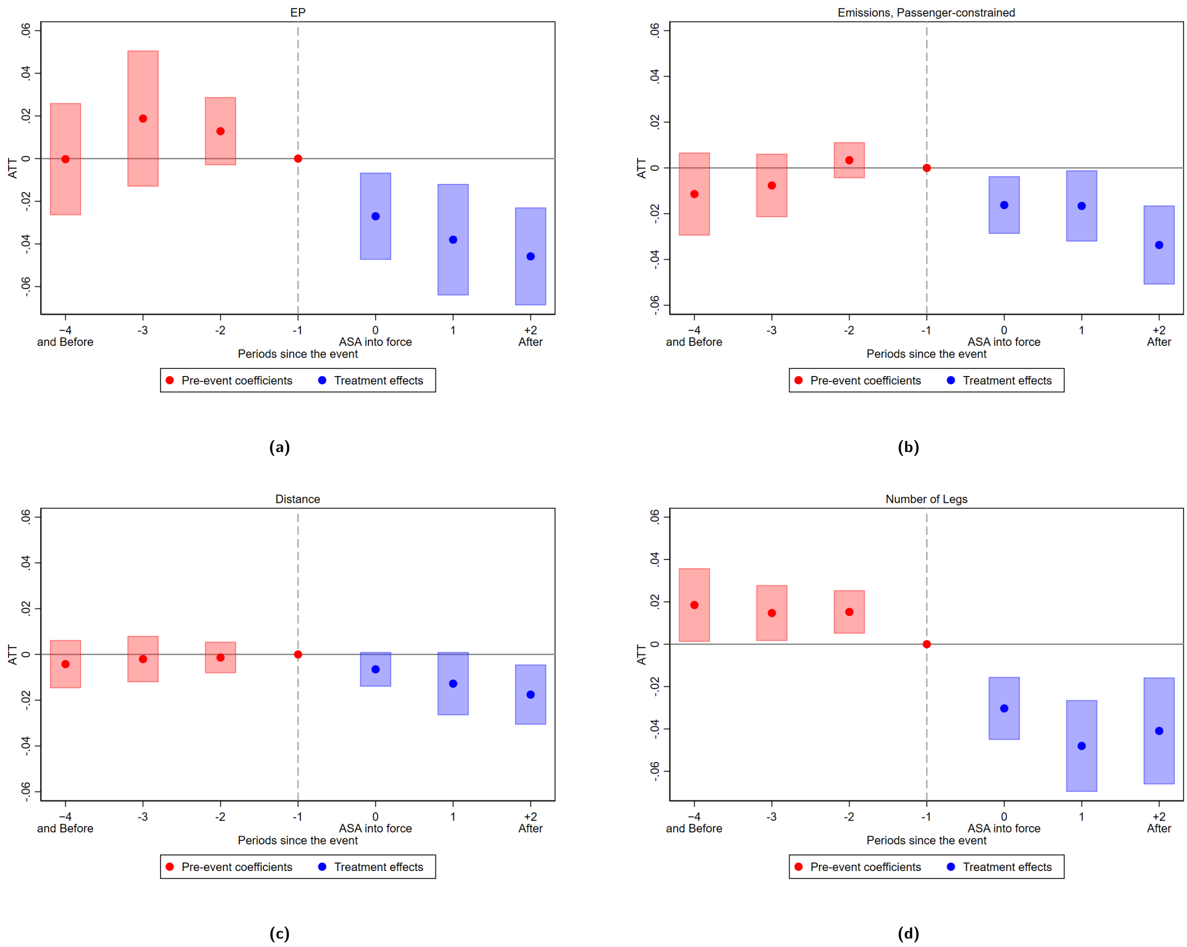

(a) (b) (c) (d)

Note: The figure reports pre- and post-treatment estimates from a PPML estimation of equation 2. The regression is estimated using a control group consisting only of never-treated observations. Cohort-time-specific treatment effects are then aggregated to obtain event-time-specific estimates, as described in section 3.2. 95% confidence intervals are reported, with standard errors clustered by country of origin and destination.

4.2. Quantification of the net effect of ASAs

The reduced-form approaches discussed above do not account for the complex interaction between airline efficiency and customer demand for air transportation services. Shorter and more comfortable direct flights enhance efficiency (technical effect), stimulate demand for international travel (scale effect), and alter passengers' choice of destinations and airlines (composition effect). The treatment of a country-pair can also affect the traffic between other country-pairs. To evaluate the broader implications of ASAs for international air transport and CO_2 emissions and to capture indirect effects, we employ a structural gravity model akin to a full endowment general equilibrium framework in the trade literature,[54] calibrated using the elasticities obtained with the TWFE reported in Table 5.[55]

This approach explicitly accounts for the complex interactions among air traffic flows, airline network optimization, and emissions, allowing us to quantify how global air traffic patterns are reshaped by ASAs. We do so by modeling bilateral passenger flows and associated emissions while imposing market-clearing conditions. We assume a fixed "endowment" of air transport capacity, such that the total productive capacity of the aviation system – proxied by the number of aircraft or available seat capacity – remains constant across simulations. Market clearing therefore occurs through adjustments in effective flying time and route efficiency. The system returns to equilibrium via implicit "price" adjustments, which reflect reductions in iceberg-type trade costs. This is an application of a standard structural gravity model for trade to the international aviation sector. In this case, the usual multilateral resistance terms have to be re-interpreted as generalized "prices", or equivalently, opportunity costs, expressed in units of flying time or travel effort. These terms reflect the relative ease or difficulty of reaching all destinations from a given origin (outward resistance) or of being reached from all origins (inward resistance).

The structure of the global ASAs network further supports this interpretation. Indeed, as of 2012, the global network of ASAs already exhibited the characteristics of a large and highly interconnected structure, in which each country was connected to all potential partners, with an average path length of approximately two intermediate connections (i.e. average path length of 1.83). Under this configuration, new ASAs primarily reduce generalized bilateral costs by shortening effective flying time or improving connectivity rather than creating new links between disconnected markets. As a result, the reallocation of passenger flows happens along the intensive margin, and the multilateral resistance terms completely capture the direct and indirect adjustments of all these relative "prices". Thus, a reduction in the resistance term, whether outward or inward, corresponds to an improvement in network accessibility, analogous to a decline in the price index of traded goods in standard gravity models. Such a structural-gravity interpretation of ASA formation is challenged only in the relatively rare case in which a previously disconnected country pair becomes directly linked by a new ASA. This constitutes a genuine extensive-margin shock that falls outside the standard structural-gravity propagation mechanism. We address this issue in the next section.

In order to calibrate the model used for the counterfactual analysis, we need to estimate the impact of ASAs on passengers and emissions separately. This is done in columns (1) and (2) of Table 5. Namley, we estimate the impact of ASAs on the total number of passengers

[54] See Anderson et al. (2018).

[55] The robustness check using the full sample is reported in Appendix Table A8.

Table 5 – The impact of ASAs on number of passengers and emissions. Period 2012-2019.

Dep Var. :	Passengers (1)	Emissions (2)
Estimator	Heterogeneity-robust DiD (never-treated)	
ASA	0.076	0.065***
	(0.051)	(0.0238)
Obs.	138,310	138,310
Estimator	Heterogeneity-robust DiD (never- & not-yet-treated)	
ASA	0.127**	0.091***
	(0.058)	(0.024)
Obs.	138,310	138,310
Estimator	Structural gravity (TWFE)	
ASA	0.086***	0.044**
	(0.023)	(0.019)
Obs.	176,638	177,638
Origin-Year FE	Yes	Yes
Destination-Year FE	Yes	Yes
Origin-Destination FE	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variables are total number of passengers and emissions. Standard errors are two-way clustered at the level of origin and destination countries * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. When the control group is restricted to never-treated country pairs, all pre-treatment observations of not-yet-treated pairs are absorbed by cohort-year fixed effects; this procedure leaves the number of observations unchanged.

and total emissions within each treated pair of countries using both heterogeneity-robust DID with two different control groups (panel a-b of Table 5), and a TWFE estimator (panel c of Table 5). Since the TWFE estimator is consistent with the theoretical framework underlying the counterfactual analysis, we employ TWFE-based elasticities in the simulations. For the counterfactual analysis that differentiates ASAs by their degree of ambition, we rely on the estimates reported in Table A6 in the Appendix.

Our counterfactual simulations compare three scenarios: (i) a world without existing ASAs (simulating a more restrictive aviation environment – the negative of the outcome can be interpreted as the average impact of the current network of agreements); (ii) a world with ASA of any type (reflecting moderate liberalization); and (iii) a Full-liberalization ASA world (simulating an open aviation market with minimal restrictions). In the first scenario, we turn off all the existing ASAs (regardless of whether they are fully liberalized or not). In the second scenario, pairs of countries without an ASA are now assumed to have a ASA of any type, while pairs of countries having an ASA (regardless of the type) do not change their agreement. In the third scenario, all country pairs are assumed to have a Full-Liberalization ASA: the global sky is fully liberalized. Each scenario captures both demand-side effects (changes in international passenger flows due to price adjustments and connectivity) and supply-side responses (optimization of airline routes and changes in the number of stopovers). A key feature of the model is that it incorporates endogenous adjustments in airfares and emissions per passenger based on network efficiency. A more detailed discussion of the counterfactual exercise is reported in Appendix section B.

As shown in Table 6, ASA liberalization increases international passenger traffic, but has different implications for total and per-passenger emissions. In the first scenario, removing all existing ASAs would reduce international passenger traffic by 2.8% and total emissions by 0.9%. Equivalently, the current network of ASAs raises total emissions by 0.9% and passenger traffic by an even larger margin, implying a 1.9% decrease in emissions per passenger relative to a world without ASAs. If all country pairs adopted an ASA of any type, international passenger traffic would increase by 2.3%, while total emissions would rise by only 1.2%. This would lead to a 1.08% decrease in emissions per passenger due to efficiency gains. A fully liberalized regime would have the strongest effects: international passengers would increase by 4.0%, total emissions would rise by 1.7%, and emissions per passenger would decrease by 2.3%. This underscores how route optimization under full liberalization can reduce aviation's environmental footprint.

Our results highlight the dual impact of ASAs: while air transport liberalization increases total emissions due to higher demand, the reorganization of flight networks and the reduction of stopovers improve efficiency, leading to lower emissions per passenger. The fully liberalized scenario generates the highest environmental efficiency gains, suggesting that optimizing routes and competing among airlines helps reduce emissions per passenger. It does so without significantly increasing emissions (+1.7% as opposed to +1.2%) while providing more benefits to passengers, as revealed by a 4.0% increase in the demand for flights. These effects are not trivial: the scale effect (the increase in air traffic) corresponds to seven months of air traffic growth, while the technical effect amounts to fifteen months of additional efficiency.

Table 6 – Full Endowment General Equilibrium Effects of Aviation Agreements on Air Pollution

	% Change in	
Scenario	# Pass. Int. Flights	CO_2 Emissions
Without existing ASA	-2.8	-0.9
With ASA (any type)	+2.3	+1.2
With fully liberalized ASA	+4.0	+1.7

Note: The table presents the full endowment general equilibrium effects of air liberalization on international flight traffic and CO_2 emissions. The counterfactual scenarios compare (i) a world without existing ASA air agreements, (ii) a world with moderate liberalization (general ASA air agreements), and (iii) a fully liberalized ASA world. While liberalization increases total emissions, emissions per passenger decline due to network efficiency gains.

4.3. The Diversion Effects and Route Centrality of ASA

The presence of an ASA connecting two locations might divert passengers (and hence revenues) from non-ASA connected pairs towards ASA-connected pairs. Indeed, the presence of a more comfortable air-route might affect the preferences of passengers and divert overall demand toward ASA-connected pairs. The General Equilibrium exercise discussed in the previous section considers the diversion effect of ASA *via* the adjustment in the country-year specific factors, i.e. average price (or Multilateral Resistance terms in the structural gravity jargon). However, to strengthen the validity of the counterfactual exercise, in this section we explicitly tests whether the signature of a new ASA diverts passengers from non-ASA to ASA connected destinations.

This is tested in Table 7 where, in addition to the dummy variable for the presence of an ASA between origin o and destination d, we include the product of the number of ASAs signed, respectively, by country o and d with third countries k.[58] This variable increases when both origin and destination have a large number of ASA agreements with third countries k, potentially diverting passengers away from the od specific route. Reassuringly, results reported in Table 7 show the absence of diversion effect.

One may also be concerned that ASAs introduce new *potential* air routes that would not otherwise exist. In this specific case, the country–year adjustment in the GE exercise may not be sufficient to capture – and therefore control for – the opening of new routes induced by ASAs. We address this concern by directly testing whether ASAs affect the structure of the air-route network.

To address this concern, we construct a measure of corridor centrality based on the structure of the global country-level air network. We first aggregate airport-to-airport great-circle distances (adjusted for Earth curvature) to the country-pair level using ticket weights, thereby obtaining a directed country network in which bilateral links are weighted by the average geographic distance between connected airport pairs. On this country-level network, we compute the distance-minimizing path between every ordered origin–destination country pair. Corridor centrality is then defined as the number of ordered country pairs whose shortest-distance path transits through a given bilateral corridor.[60]

This shortest-path corridor centrality captures the structural importance of a country pair within the global routing architecture: a corridor is central if it lies on many optimal (distance-minimizing) country-to-country itineraries connecting the rest of the network. Importantly, the global air network forms essentially a single giant connected component, so new ASAs do not create connectivity where none previously existed. Rather, they modify relative generalized costs within an already connected system.

If ASAs disproportionately benefited already pivotal backbone corridors, we would expect a significantly stronger effect on highly central routes. Columns (3) and (4) of Table 7 show that this is not the case. The interaction between ASA adoption and high centrality is negative and statistically insignificant, indicating that the passenger gains from ASAs are not concentrated on structurally dominant corridors. Combined with the absence of diversion effects reported in columns (1) and (2), these results support the interpretation that ASAs primarily operate through bilateral cost and efficiency channels embedded in multilateral resistance, rather than through systematic demand reallocation or large-scale reconfiguration of the network structure.

Overall, the absence of both diversion effects and changes in the network structure of air routes reassures us that omitted-variable bias is unlikely to affect the estimated parameters used in the GE exercise, thereby strengthening the validity of the counterfactual results.

[58]To prevent any scale effect in the constriction of this variable, we divide the number of ASAs of country o and d with third countries k by the maximum possible number of ASAs

[60]We compute the distance-minimizing path between every ordered origin–destination pair using Dijkstra's algorithm Dijkstra (1959).

Table 7 – The diversion effect of ASAs and route centrality.

Dep Var.:	Revenues	Passengers	Passengers High Centrality Corridor $>$ Mean	 $>$ Median
	(1)	(2)	(3)	(4)
ASA$_{odt}$	0.031*	0.086***	0.104***	0.118***
	(0.016)	(0.023)	(0.029)	(0.041)
$\sum_k ASA_{okt} \times \sum_k ASA_{dkt}$	-0.008	-0.002		
	(0.011)	(0.013)		
ASA$_{odt}\times$ High Centrality$_{od}^{2012}$			-0.072	-0.050
			(0.052)	(0.040)
Obs.	176,638	176,638	176,638	176,638
Estimator	Structural gravity (TWFE)			
Origin-Year FE	Yes	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes	Yes
Origin-Destination FE	Yes	Yes	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. Columns (1)–(2) report the baseline diversion specification. Columns (3)–(4) allow for heterogeneous effects by interacting ASA$_{odt}$ with an indicator for high corridor centrality. Corridor centrality is defined as the number of ordered country pairs whose shortest-distance path (computed on the country-level network using great-circle distances) transits through the bilateral corridor. "High Centrality" corresponds to corridors above the sample mean (column 3) or median (column 4). Standard errors are two-way clustered at the level of origin and destination countries. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5. Conclusion

This paper combines a reduced form econometric approach and a counterfactual analysis based on a structural gravity model accounting for all indirect effects of ASAs to study the environmental impact of Air Services Agreements. To do so, we use granular data on passenger tickets from 2012 to 2019 from SABRE, as well as information on aircraft emissions, flight routes, number of legs, airport location, and the presence of Air Service Agreements. Our reduced-form econometric analysis suggests that ASAs reduce the average amount of emissions per passenger by 3.9% in our preferred specification. One half of this impact is driven by the increase in efficiency in international routes, as revealed by a decrease in the average distance and number of legs. Composition effects, where passengers adjust their itineraries and arbitrage between airlines, explain the remaining part of the overall effect. Interestingly, the degree of liberalization matters, and Full-liberalization ASAs are driving the observed reduction in emissions per passenger. The inclusion of a comprehensive set of fixed effects and the absence of pre-trend (suggested by an event study approach) allow us to conclude that ASAs have a *causal* impact on international flight distance, stopovers, and consequently, emissions per passenger.

The counterfactual analysis suggests that the existing network of ASAs ensures a -1.6 percent reduction in emissions per passenger. A further liberalization of the air services, in which all country pairs would be linked by ambitious ASAs (full liberalization), would imply a 2.3% reduction in emissions per passenger and a 4.0% increase in the number of passengers; hence, a 1.7% increase in total emissions. The main message of the counterfactual analysis is therefore that the scale effect (increase in the total number of passengers) of ASAs consistently outweighs the efficiency effect (reduction in average distance).

References

Anderson, J. E., Larch, M., and Yotov, Y. V. (2018). GEPPML: General equilibrium analysis with PPML. *The World Economy*, 41(10):2750–2782.

Arlines.org (2024). World Airlines Traffic and Capacity. https://airlines.org/dataset/world.airlines-traffic-and-capacity/. Accessed 5 March 2024.

Bahar, D., Choudhury, P., Kim, D. Y., and Koo, W. (2023). Innovation on wings: Nonstop flight and firm innovation in the global context. *Management Science*, 69(3):6202–6223.

Bekkers, E. (2019). The welfare effects of trade policy experiments in quantitative trade models: The role of solution methods and baseline calibration. WTO Staff Working Papers ERSD-2019-02, World Trade Organization (WTO), Economic Research and Statistics Division.

Blonigen, B. A. and Cristea, A. D. (2015). Air service and urban growth: Evidence from a quasi-natural policy experiment. *Journal of Urban Economics*, 86:128–146.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6):3253–3285.

Campante, F. and Yanagizawa-Drott, D. (2018). Long-range growth: economic development in the global network of air links. *The Quarterly Journal of Economics*, 133(3):1395–1458.

Copeland, B. R. and Taylor, M. S. (1993). North-south trade and the environment. *Quarterly Journal of Economics*, 109(3):755–787.

Cristea, A. D. (2011). Buyer-seller relationships in international trade: Evidence from us states' exports and business-class travel. *Journal of international economics*, 84(2):207–220.

De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

De Chaisemartin, C. and d'Haultfoeuille, X. (2023). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey. *The Econometrics Journal*, 26(3):1–30.

Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271.

Edgar (2024). Emissions Database for Global Atmospheric Research. https://edgar.jrc.ec.europa.eu/emissions_data_and_maps. Accessed 5 March 2024.

Fageda, X. and Teixidó, J. J. (2022). Pricing carbon in the aviation sector: Evidence from the european emissions trading system. *Journal of Environmental Economics and Management*, 111:102591.

Gabrielli, L., Deutschmann, E., Natale, F., Recchi, E., and Vespe, M. (2019). Dissecting global air traffic data to discern different types and trends of transnational human mobility. *EPJ Data Science*, 8(1):1–24.

Grossman, G. M. and Krueger, A. B. (1993). 2 environmental impacts of a north american free trade agreement. *The Mexico-US Free Trade Agreement*, 11(2):13.

Ho, Chun-Yu and Peng, Tingting and Takayama, Haruka and Xu, Li (2024). Air Connectivity and International Travel:Evidence from Cross-border Card Payments. https:/hatakayama.github.io/paper_web/Air_Connectivity_and_Cross_border_Travel.pdf.

Hovhannisyan, N. and Keller, W. (2015). International business travel: an engine of innovation? *Journal of Economic Growth*, 20:75–104.

IATA (2019). Annual Review. https://www.iata.org/en/publications/annual-review/. Accessed 5 March 2024.

ICAO Carbon Emissions Calculator (2024). https://www.icao.int/environmental-protection/Carbonoffset/Pages/default.aspx. Accessed 5 March 2024.

Larch, M. and Wanner, J. (2024). The consequences of non-participation in the paris agreement. *European Economic Review*, 163:104699.

Lee, D., Fahey, D., Skowron, A., Allen, M., Burkhardt, U., Chen, Q., Doherty, S., Freeman, S., Forster, P., Fuglestvedt, J., Gettelman, A., De León, R., Lim, L., Lund, M., Millar, R., Owen, B., Penner, J., Pitari, G., Prather, M., Sausen, R., and Wilcox, L. (2021). The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. *Atmospheric Environment*, 244:117834.

Morales-Arilla, J. and Bustos, S. (2024). Direct flights, trade and specialization. *mimeo, Harvard*.

Myflight (2024a). https://co2.myclimate.org/fr/flight_calculators. Accessed 5 March 2024.

Myflight (2024b). https://www.myclimate.org/fr-ch/sinformer/notre-profil/telechargements/calculateur-de-vol/. Accessed 5 March 2024.

Nagengast, A. J. and Yotov, Y. V. (2025). Staggered difference-in-differences in gravity settings: Revisiting the effects of trade agreements. *American Economic Journal: Applied Economics*, 17(1):271–296.

Piermartini, R. and Rousová, L. (2013). The sky is not flat: how discriminatory is the access to international air services? *American Economic Journal: Economic Policy*, 5(3):287–319.

Roth, J., Sant'Anna, P. H., Bilinski, A., and Poe, J. (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.

Santos-Silva, J. M. C. and Tenreyro, S. (2006). The Log of Gravity. *The Review of Economics and Statistics*, 88(4):641–658.

Seymour, K., Held, M., Georges, G., and Boulouchos, K. (2020). Fuel estimation in air transportation: Modeling global fuel consumption for commercial aviation. *Transportation Research Part D: Transport and Environment*, 88:1025–28.

Shapiro, J. S. (2025). Institutions, comparative advantage, and the environment. *The Review of Economic Studies*, page rdaf012.

Söderlund, B. (2023). The importance of business travel for trade: Evidence from the liberalization of the soviet airspace. *Journal of International Economics*, 145:103812.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199. Themed Issue: Treatment Effect 1.

Wang, Z., Wang, F., and Zhou, Z. (2021). All roads lead to rome: global air connectivity and bilateral trade. *Available at SSRN 3884374*.

Wooldridge, J. M. (2021). Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators. Working paper, SSRN.

Wooldridge, J. M. (2023). Simple approaches to nonlinear difference-in-differences with panel data. *The Econometrics Journal*, 26(3):C31–C66.

Yan, Z., Zhang, J., Wang, Z., and Du, Z. (2025). Does international coordinated industrial policy stimulate regional low-carbon aviation development? evidence from corsia and the apec developing economies. *Energy Economics*, page 108220.

A. Additional figures and tables

Table A1 – Values to calculate CO2 EP

	Short haul Standard < 1500 Km	Long haul Standard > 2500 Km
S: Seats	157.86	302.58
PLF: Passenger load factor	0.796	0.82
CF: Cargo Factor	0.26	0.26
CW: Cabin class weighting factor		
Economic	1	1
Economic premium	1	1.5
Business	1.5	4
First Class	1.5	5
EF: CO2 emission factor	3.16	3.16
Fluel parameters		
a	0.000007	0.00029
b	2.588	3.475
c	1260.608	3259.691

Source: data provided by Myflight (2024b)

Table A2 – Summary statistics within origin–destination pairs

Variable	N	Mean	SD	Min	Max
County level – within origin–destination pairs					
Tot. Emissions (Gg CO_2 eq.)	176638	28,71	728,36	0,000031	113114,51
Passengers (Mio)	176638	0,139	4,92	0,000001	731,62
EP (Kg CO_2 eq.)	176638	653,66	476,57	24,85	6763,32
EPKM (Kg CO_2 eq.)	176638	0,071	0,024	0,0049	0,86
Distance (Km)	176638	8159,6	4788,77	29	31707
# legs	176638	2,5	0,7	1	4
Firm level – within origin–destination pairs					
Tot. Emissions (Gg CO_2 eq.)	2765394	1,83	65,14	0,00003	21265,97
Passengers (Mio)	2765394	0,01	0,45	0,000001	158,74
EP (Kg CO_2 eq.)	2765394	743,8	647,9	23,9	13548,7
EPKM (Kg CO_2 eq.)	2765394	0,08	0,04	0,003	3,3
Distance (Km)	2765394	8736,8	5063,4	9,0	42957,0
# legs	2765394	2,7	0,7	1	4

Note: Statistics are calculated using the main county-level and firm-level samples employed in the baseline regressions.

Table A3 – The impact of ASAs on CO2 emissions, average distance and number of legs. Robustness check using the full sample.

Dep Var. :	Tot. emissions (1)	EP (2)	EPKM (3)	Distance (4)	# legs (5)
Estimator	Heterogeneity-robust DiD (never treated)				
ASA	-0.021***	-0.029***	-0.033***	-0.014**	-0.030***
	(0.0049)	(0.0100)	(0.009)	(0.006)	(0.008)
Obs.	138,567	138,483	138,483	138,483	138,483
Estimator	Heterogeneity-robust DiD (never- & not-yet-treated)				
ASA	-0.021***	-0.039***	-0.044***	-0.013**	-0.041***
	(0.006)	(0.0113)	(0.0107)	(0.005)	(0.010)
Obs.	138,567	138,483	138,483	138,483	138,483
Estimator	Structural gravity (TWFE)				
ASA	-0.015	-0.028**	-0.039***	-0.011*	-0.044***
	(0.010)	(0.0014)	(0.0014)	(0.006)	(0.011)
Obs.	177,020	177,020	177,020	177,020	177,020
Origin-Year FE	Yes	Yes	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes	Yes	Yes
Origin-Destination FE	Yes	Yes	Yes	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variables are defined as follows: column (1) reports total emissions with log passenger numbers included as an offset, which constrains the elasticity of emissions with respect to passengers to one ("passenger-constrained"); column (2) reports emissions per passenger (EP); column (3) reports emissions per passenger-kilometer (EPKM); column (4) reports the average distance traveled between origin and destination; and column (5) reports the average number of flight segments. Standard errors are two-way clustered at the level of origin and destination countries. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. When the control group is restricted to never-treated country pairs, all pre-treatment observations of not-yet-treated pairs are absorbed by cohort-year fixed effects; this procedure leaves the number of observations unchanged.

Table A4 – The impact of ASAs on revenues and average price. Period 2012-2019.

Dep Var. :	Revenues	Average Price	Average Price (per Km)	Average Price (Dist. constrained)
	(1)	(2)	(3)	(4)
Estimator	Heterogeneity-robust DiD (never-treated)			
ASA	0.035	-0.015	0.0037	0.0007
	(0.027)	(0.012)	(0.015)	(0.014)
Obs.	138,310	138,310	138,310	138,310
Estimator	Heterogeneity-robust DiD (never- & not-yet-treated)			
ASA	0.068**	0.001	0.046***	0.017
	(0.031)	(0.014)	(0.015)	(0.015)
Obs.	138,310	138,310	138,310	138,310
Estimator	Structural gravity (TWFE)			
ASA	0.033**	0.012	0.044***	0.026*
	(0.016)	(0.013)	(0.013)	(0.015)
Obs.	176,638	176,638	176,638	176,638
Origin-Year FE	Yes	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes	Yes
Origin-Destination FE	Yes	Yes	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variable in column (1) is total revenues. Columns (2)-(4) show results on average price, average price per Km and distance-constrained average price. Standard errors are two-way clustered at the level of origin and destination countries * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. When the control group is restricted to never-treated country pairs, all pre-treatment observations of not-yet-treated pairs are absorbed by cohort-year fixed effects; this procedure leaves the number of observations unchanged.

Table A5 – The impact of ASAs on emissions per passenger by type of ASA. Robustness check using the full sample.

Dep Var.	Emissions per passenger (EP)		
	(1)	(2)	(3)
ASA Full lib.	-0.056***	-0.063***	-0.054***
	(0.012)	(0.0137)	(0.015)
ASA Not-Full Lib.	0.005	-0.012	0.004
	(0.012)	(0.012)	(0.014)
Obs.	138,657	138,657	177,020
Estimator	Heterogeneity-robust DiD never-treated	Heterogeneity-robust DiD never- & not-yet-treated	Structural gravity (TWFE)
Origin-Year FE	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes
Origin-Destination FE	Yes	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variable is emissions per passenger (EP). Standard errors are two-way clustered at the level of origin and destination countries * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. When the control group is restricted to never-treated country pairs, all pre-treatment observations of not-yet-treated pairs are absorbed by cohort-year fixed effects; this procedure leaves the number of observations unchanged.

Table A6 – The impact of ASAs on the number of passengers and emissions by type of ASA.

Dep Var.	Passengers (1)	Emissions (2)
ASA Full lib.	0.105***	0.052**
	(0.024)	(0.021)
ASA Not-Full Lib.	0.047*	0.036*
	(0.025)	(0.022)
Obs.	176,638	176,638
Estimator	Structural gravity (TWFE)	
Origin-Year FE	Yes	Yes
Destination-Year FE	Yes	Yes
Origin-Destination FE	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. Standard errors are two-way clustered at the level of origin and destination countries * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7 – The impact of ASAs on emissions per passenger. Robustness checks using different samples.

Dep Var.	EP			
	All	IATA only	Main sample (Distance > 1500 Km)	Main sample (Excluding national flight)
	(1)	(2)	(3)	(4)
ASA	-0.0147**	-0.0143**	-0.0142**	-0.0143**
	(0.005)	(0.006)	(0.007)	(0.006)
Obs.	1,734,907	1,563,619	1,520,435	1,534,172
Estimator	Heterogeneity-robust DiD never- & not-yet-treated			
Origin-Year FE	Yes	Yes	Yes	Yes
Destination-Year FE	Yes	Yes	Yes	Yes
Origin-Destination Airline FE	Yes	Yes	Yes	Yes
Airline-Year FE	Yes	Yes	Yes	Yes

Note: In columns (1) we show results using the full sample (including also non-commercial flights and observations below bottom-1 percentile). In columns (2) we show results using commercial flights and include bottom-1 percentile observations. Columns (3) and (4) show results using respectively long distance flights (above 1500 Km) and excluding domestic flights. Estimates are obtained using PPML over the period 2012–2019. The dependent variable is emissions per passenger (EP). Standard errors are three-way clustered at the origin, destination, and airline-year levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8 – The impact of ASAs on the number of passengers and emissions. Robustness check using the full sample.

Dep Var. :	Passengers (1)	Emissions (2)
Estimator	Heterogeneity-robust DiD (never-treated)	
ASA	0.059	0.057**
	(0.055)	(0.025)
Obs.	138,657	138,657
Estimator	Heterogeneity-robust DiD (never- & not-yet-treated)	
ASA	0.112*	0.085***
	(0.063)	(0.027)
Obs.	138,657	138,657
Estimator	Structural gravity (TWFE)	
ASA	0.101***	0.056*
	(0.029)	(0.024)
Obs.	177,020	177,020
Origin-Year FE	Yes	Yes
Destination-Year FE	Yes	Yes
Origin-Destination FE	Yes	Yes

Note: Estimates are obtained using PPML over the period 2012–2019. The dependent variables are total number of passengers and emissions. Standard errors are two-way clustered at the level of origin and destination countries * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. When the control group is restricted to never-treated country pairs, all pre-treatment observations of not-yet-treated pairs are absorbed by cohort-year fixed effects; this procedure leaves the number of observations unchanged.

B. Full Endowment GE Framework for ASAs

In order to quantify the impact of Air Service Agreements (ASAs) on international aviation and CO_2 emissions, we use a structural gravity model akin to a full-endowment general equilibrium framework in the trade literature (Anderson et al., 2018), here applied to the aviation sector. This approach provides a theory-consistent method to estimate how changes in aviation policy affect passenger flows, airfares, and emissions across interconnected markets. In the context of environmental policy, structural gravity frameworks have also been extended to model the impact of carbon emissions and leakage from international policy cooperation – see Larch and Wanner (2024) among others.[72] In the context of air transport, this approach allows us to evaluate counterfactual scenarios while ensuring internal consistency across all bilateral aviation routes.

The model assumes that the number of passengers traveling from country o to country d in year t is determined by:

$$\text{Passengers}_{odt} = \frac{E_{ot} \cdot C_{dt}}{\Phi_{ot}\Psi_{dt}} \cdot \tau_{od}^{-\beta}, \tag{3}$$

where E_{ot} represents the total number of outbound passengers (or available flight capacity) from country o at time t, C_{dt} captures the total demand for international flights to country d at time t, and τ_{odt} denotes the bilateral time-varying air transport costs, which are affected by ASAs. The terms Φ_{ot} and Ψ_{dt} represent, respectively, the outward and inward multilateral resistance terms, which account for the fact that a change in any bilateral flight route affects all other routes. Following Anderson et al. (2018), we estimate this equation using Poisson Pseudo-Maximum Likelihood (PPML), which corrects for heteroskedasticity (Santos-Silva and Tenreyro, 2006) and, thanks to its additivity properties, allows the model to be solved in a straightforward way without the need for nonlinear solvers.

The process involves iteratively updating the equilibrium values of bilateral passenger flows and the multilateral resistance terms. After estimating the baseline gravity specification via PPML, we recover the fixed effects to obtain the multilateral resistance terms. We then simulate counterfactual policy shocks, such as the elimination or liberalization of ASAs, by adjusting bilateral aviation costs, τ_{odt}. We solve the system of equilibrium conditions once again in such a way that the revised bilateral and total passenger flows are consistent with each other, thereby achieving market clearing. This iterative process continues until the changes in the endogenous variables converge to a stable equilibrium that captures both the direct effects of the policy shock on bilateral traffic and the indirect effects that are transmitted through the network of multilateral resistances.[73]

We consider three counterfactual scenarios:

1. A world without existing PTA air agreements: the negatives of the simulated effects represent the direct and indirect impacts of the existing network of agreements;

[72] Larch and Wanner (2024) adopt a full-endowment general equilibrium framework to quantify the consequences of non-participation in the Paris Agreement

[73] This iterative procedure allows us to obtain stable general equilibrium indices that are robust to the specification of the trade cost vector; see Bekkers (2019) for further discussion.

2. A world with general ASAs (no matter their depth): this simulates a situation where all country-pairs without an ASA sign one, accounting for the average effects of ASAs;
3. A fully liberalized ASA world: this simulates a situation where all country-pairs shift to a Fully Liberalized ASA, starting from a situation with or without an ASA.

These simulations capture both the demand-side effects of increased international connectivity and the supply-side adjustments in airline route optimization and pricing. By iterating over new price levels and traffic volumes, we ensure that the multilateral resistance terms adjust accordingly (Anderson et al., 2018).

As shown in Table 6 of the main text, our results highlight a trade-off: while total emissions increase due to higher flight demand, emissions per passenger decline as ASAs facilitate more direct flights and reduce unnecessary stopovers, thereby improving network efficiency. By leveraging this approach, we capture not only the direct effects of ASAs on passenger flows but also the broader network effects that shape aviation markets and emissions. Our results reinforce the idea that aviation liberalization can improve environmental efficiency by reducing per-passenger emissions, even as total emissions increase due to rising demand.