Transportation Research Part E 208 (2026) 104688

Contents lists available at ScienceDirect

Transportation Research Part E

journal homepage: www.elsevier.com/locate/tre

The structure of global delay propagation in air transport

Josu Blanco, Antònia Tugores, José J. Ramasco, Massimiliano Zanin *

Instituto de Física Interdisciplinar y Sistemas Complejos IFISC (CSIC-UIB), Palma de Mallorca, Spain

ARTICLE INFO

Keywords:
Air transport
Delay propagation
Granger causality

ABSTRACT

The modelling and understanding of how flight delays propagate between airports of a country or region is a major topic of research in air transport that has been tackled through different techniques in the last decade. Much less attention has been devoted to the large-scale structure of the propagation, i.e. if and how delays can propagate between regional networks and continents. By leveraging on two complementary analysis approaches, we show how such propagation takes place across the main world regions, describing how it is modulated by seasons, the number of flights connecting them and their relative distance. We further propose a methodology to detect which airports act as gateways for global-scale propagation, and discuss the operational applications of these findings.

1. Introduction

Describing and modelling delay propagation patterns in air transport is a topic that has long attracted the interest of the research community (Li et al., 2024) due to the large impact that these have - both in terms of direct and indirect costs for airports, airlines, passengers (Cook and Tanner, 2011; Song et al., 2020), and the environment (Carlier et al., 2007). Several proposals rely on simulations that model aircraft and passenger movements through the system, including operational rules such as requiring one flight to wait for another that is delayed (Hansen, 2002; Fleurquin et al., 2013; Pyrgiotis et al., 2013; Baspinar and Koyuncu, 2016; Li et al., 2020). Building such models is nevertheless a complex task, as many relevant factors may not be known - for instance, internal rules of airlines to decide whether to wait for connecting passengers, or the sequence of flights operated by the same aircraft and crews. As a consequence, an alternative approach has been proposed, based on extracting instances of propagation from real time series representing the dynamics of airports (Zanin, 2015; Zanin et al., 2017; Du et al., 2018; Pastorino and Zanin, 2021; Jia et al., 2022). In short, metrics and tests, like Granger Causality (GC) (Granger, 1969), are used to detect whether the evolution of delays at one airport drives the evolution of delays at a second one; and the complete propagation patterns, i.e. the instances that are detected between all possible pairs of airports, are then represented as functional complex networks (Strogatz, 2001; Newman, 2003). While this approach presents the advantage of relying only on delay time series, which are easy to obtain, it also entails important challenges. Among others, propagation instances can reliably be detected only in long time series, when these are properly detrended, and when the intensity of the propagation is large enough (Acharya et al., 2024; Olivares et al., 2025).

Long-distance delay propagation, such as between intercontinental airports, underscores the importance of these limitations related to the duration and intensity of the time series used for analysis. On the one hand, one may intuitively expect those propagation instances to exist: an intercontinental flight may arrive late due to a delay at departure, and in turn, regional flights have to be delayed to allow connecting passengers to continue their trip. Yet, on the other hand, those flights are few and of low frequency. To illustrate, only tens of flights per day connect cities in Europe and US; when compared to the thousands of flights operating in each

* Corresponding author.
E-mail address: massimiliano.zanin@gmail.com (M. Zanin).

https://doi.org/10.1016/j.tre.2026.104688
Received 23 May 2025; Received in revised form 9 January 2026; Accepted 9 January 2026
Available online 27 January 2026

of these networks, the contribution of intercontinental ones to the average delay observed at each individual airport looks minor at best. In other words, the signal-to-noise ratio is expected to be very weak. It is therefore not surprising that, while many studies have been conducted regarding the propagation of delays within e.g. Europe (Zanin, 2015; Pastorino and Zanin, 2021; Zanin, 2021; Belkoura and Zanin, 2016), the U.S. (Mazzarisi et al., 2020; Bombelli and Sallan, 2023) or China (Zanin et al., 2017; Jia et al., 2022; Xu and Zhang, 2022; Zhang et al., 2022; Pastorino and Zanin, 2023), to the best of our knowledge, none has reported results at an interregional level.

In this contribution, we characterise the structure of air traffic delay propagation between continents, and develop a methodology capable of detecting the weak signals characteristic of these long-range interactions. Applying this methodology to a comprehensive global data set for the period 2015–2018, we map the global delay network and identify its key drivers. The detailed contributions can be summarised as:

- We develop and validate two complementary frameworks, a microscale and a macroscale approach, to detect long-range flight delay propagation signals. Validation is achieved using official regional data sets from EUROCONTROL (EUROCONTROL, 2025a) and the U.S. Bureau of Transportation Statistics (BTS) (Bureau of Transportation Statistics, U. S. Department of Transportation, 2025). The main analysis is performed using a global FlightRadar24 (FR24) (FlightRadar24, 2025b) data set.
- We provide, to the best of our knowledge, the first empirical evidence that flight delays systematically propagate across continents, with the resulting global network of delay propagation mapped explicitly.
- We analyse the structure of these global patterns, showing modulation by seasonality, interregional traffic volume and geographical distance.
- We introduce a recursive feature elimination methodology to identify and rank the key airports functioning as sources of intercontinental delay propagation.

The remainder of the work is organised as follows. Firstly, Section 2 presents a review of the state of the art on delay propagation analysis. Secondly, Section 3 introduces the data and methods underpinning our study, including a description of the data sources in Section 3.1, the pre-processing steps in Section 3.2, the Granger Causality test in Section 3.3, and a comparison of the data sets in Section 3.4. We then present the two proposed frameworks for assessing long-range delay propagation in Section 4. Next, the results of the global analysis are detailed, focusing on the large-scale propagation patterns in Section 5 and the role of individual airports in this process in Section 6. Finally, Section 7 discusses our findings, their operational implications, limitations, and future paths for research.

2. State of the art in delay propagation analysis

The global air traffic network mainly follows a hub-and-spoke topology (Bryan and O'kelly, 1999; Barthélemy, 2011), where central hub airports serve as key transfer points connecting multiple smaller airports, i.e. the spokes. This structure enhances cost efficiency by consolidating traffic through major hubs, minimising the need for direct connections between individual locations. Long-haul flights link these hubs, forming a well-interconnected network. If airports are considered as nodes and flights as edges of a network, previous studies have shown that the air transportation network exhibits small-world properties—characterised by short average path lengths between airports—alongside a scale-free topology (Barabási and Bonabeau, 2003), where the probability of a node having a certain number of connections follows a long-tailed distribution. In such networks, most nodes have relatively few connections, while a small number of highly connected hubs dominate the system (Zanin and Lillo, 2013; Cook et al., 2015). Similar structures are observed in both natural (Ravasz and Barabási, 2003) and human-made (Albert et al., 1999) systems. Notably, this contrasts with other transportation networks, such as road (Xie and Levinson, 2007) and rail (Lu et al., 2004) systems, which are more geographically constrained and generally lack such highly connected hubs.

A key consequence of this topology is the air transport network's robustness to random failures, but high vulnerability to targeted disruptions. Since most airports have low connectivity, removing a randomly chosen airport typically has little global impact (Albert et al., 2000). However, the network is especially susceptible to disruptions at major hubs, where failures can significantly impair the overall connectivity (Albert et al., 2000; Latora and Marchiori, 2003). From a delay propagation viewpoint, this implies a high sensitivity to feeder delays i.e. delays affecting flights from spokes into hubs (Ruehle et al., 2006). The sensitivity arises not because individual spokes are critical, but because hubs act as aggregators where the cumulative impact of many minor feeder delays can saturate operational capacity. The hub then amplifies these consolidated delays, propagating them system-wide through its high-volume connections. Such delays can then cascade and propagate, significantly impacting the network status.

Understanding the air traffic network as a complex system provides valuable insights into its resilience and vulnerabilities. In particular, the network's topology is essential for analysing failure cascades (Cook et al., 2015). However, resilience in this context extends beyond topology to include operational factors (Wandelt et al., 2023). Airlines incorporate buffer times and recovery procedures to mitigate delays, effectively engineering resilience (Brueckner et al., 2021). For example, schedule buffers help to ensure that late-arriving aircraft can still depart on time, thereby limiting the propagation of delays. Consequently, arrival delays tend to be shorter than departure delays (Zheng et al., 2021). Long-term adaptations, such as re-routing traffic or modifying schedules, demonstrate adaptive capacity. Research indicates that the network exhibits a 24-hour reset cycle, with the overnight decrease in operations playing a crucial role in restoring normal dynamics (Gopalakrishnan and Balakrishnan, 2021; Zanin and Lillo, 2013).

Early foundational studies of airline scheduling introduced the concept of the delay multiplier, quantifying how an initial (or primary) delay propagates through subsequent flights (Beatty et al., 1999). Delays spread primarily through crew and aircraft connections, with later research works estimating that each minute of initial delay generates, on average, an additional 1.5 to 2 minutes of

downstream delay (Boswell and Evans, 1997). Subsequent works incorporated schedule buffers, employing methods such as survival analysis and Bayesian modelling to better capture how turnaround times either mitigate or amplify delays (AhmadBeygi et al., 2008; Kim and Bae, 2021); and further analysed the nonlinear nature of such delay multiplication (Belkoura et al., 2017). In Europe, where airspace is more congested and operations span multiple national jurisdictions, network-wide delay analyses have been led by organisations such as EUROCONTROL (EUROCONTROL, 2025b). A particularly significant statistic is the high proportion of reactionary delays, i.e. delays caused by late inbound aircraft, which account for approximately 40% of total delay minutes in the European system (EUROCONTROL, 2023). This explains why adverse events, such as industrial actions or airport closures, generate widespread ripple effects across the network (Zanin and Lillo, 2013). Additionally, certain regions experience frequent weather-related delays, further compounding network disruptions (Wong and Tsai, 2012).

Granger Causality (GC) has become a valuable tool for detecting and quantifying air traffic delay propagation. Originally developed within econometrics (Granger, 1969), it has later been applied across diverse fields, including climate science (Silva et al., 2021; McGraw and Barnes, 2018), neuroscience (Seth et al., 2015; Ding et al., 2006), and financial markets (Hiemstra and Jones, 1994; Ajayi et al., 1998; Vỳrost et al., 2015). Its first application to air traffic networks appeared in (Zanin, 2015), leading to numerous studies exploring its potential. Research on European air traffic (Zanin, 2015; Pastorino and Zanin, 2021; Zanin, 2021; Belkoura and Zanin, 2016) has demonstrated its capability to identify temporal dependencies and assess the influence of network topology on delay dynamics. Similarly, studies on China's air transport system (Zanin et al., 2017; Du et al., 2018; Xu and Zhang, 2022; Zhang et al., 2022; Pastorino and Zanin, 2023) and the U.S. (Bombelli and Sallan, 2023) have highlighted the role of smaller airports and operational factors, such as aircraft type, in delay transmission. Recent advancements have refined Granger Causality to capture nonlinear dependencies in delay propagation patterns (Jia et al., 2022; Feng et al., 2024). Applications to extreme delay events (Mazzarisi et al., 2020) indicate that, while large hubs shape the average propagation dynamics, smaller airports can exert a disproportionate influence on severe disruptions due to the hub-and-spoke topology. For the sake of synthesis, some key previous works analysing delay propagation through GC and similar metrics are reported in Table 1.

Granger Causality has, of course, not been the only approach to study delay propagation. Numerical modelling and simulation-based approaches provide alternative frameworks, complementing data-driven causality analyses. Agent-based models (Fleurquin et al., 2013) have been employed to simulate the cascading effects of delays under various traffic conditions and control strategies, providing insights into systemic vulnerabilities and proposing potential mitigation measures. Queuing models have also been widely used to capture congestion effects at individual airports and their impact on network-wide performance, with frameworks like the Approximate Network Delays (AND) model being employed to quantify how an initial delay at one airport propagates throughout the network (Hansen, 2002; Pyrgiotis et al., 2013). A complementary perspective treats delay spreading as an epidemic-like process, where a heavily delayed airport is considered "infected" and can "transmit" delays to its destinations. This analogy has led to the application of Susceptible-Infected-Recovered (SIR) models at the airport level to simulate how delays propagate over time (Baspinar and Koyuncu, 2016; Li et al., 2020).

Recent advancements in machine learning have significantly enhanced predictive models, enabling real-time estimation of systematic delay patterns and their dependence on external factors such as weather conditions and operational constraints. Incorporating upstream delay information, such as delays due to late-arriving aircraft, has been shown to substantially improve prediction accuracy (Chen and Li, 2019). Beyond conventional machine learning approaches, more advanced Deep Learning techniques have also been explored, including Multi-Layer Perceptron models (MLP) (Chen et al., 2024), graph convolutional Long Short-Term Memory (LSTM) networks (Li et al., 2023; Kang et al., 2023) and attention-based spatio-temporal models (Tan et al., 2022). More recent works have proposed improving the latter by integrating heterogeneous operational data with coupled attention to improve prediction (Wu et al., 2023; Li et al., 2024). Incorporating the flight network topology directly into predictive models has enabled the development of network-informed deep learning frameworks that can identify the airports most responsible for delay propagation; by using learned attention weights, these models improve short-term delay forecasting while offering a more robust, data-driven understanding of air traffic disruptions (Kang et al., 2023).

While most studies have focused on specific regions, such as individual countries or continents, the global structure of air transportation networks has also been examined. As seen before, these networks exhibit scale-free and small-world properties, with community structures playing a significant role in shaping network dynamics and geometry (Guimera et al., 2005). Notably, the most highly connected cities are not always the most central in terms of network control, emphasising the need to account for geopolitical and structural factors when analysing global air traffic patterns. These findings underscore the importance of integrating both local and global perspectives, as regional characteristics, policy regulations and operational constraints collectively influence delay propagation and overall system resilience (Zanin and Lillo, 2013; Guimera et al., 2005).

Notably, the two aforementioned trends in the literature (i.e. delay propagation analysis and the global spatial nature of air transport) have not intersected; that is, to the best of our knowledge, a systematic analysis of the propagation of delays at the global inter-continental level has not yet been undertaken. This is illustrated in Table 1, where it can be appreciated that causality tests have always been applied to individual regions (e.g. Europe or China), but not to combinations thereof. As previously introduced, addressing this gap requires methodological innovation, as long-range propagation signals are expected to be significantly weaker than their regional counterparts, buried in the noise of thousands of local flights. This contribution fills this gap by means of the two frameworks that will be presented in Section 4, and represents the first analysis of the propagation between (as opposed to within) regions.

Table 1
List of relevant papers analysing delay propagation through causality metrics, as of November 2025. GC: Granger Causality; TE: Transfer Entropy; CCM: Convergent Cross Mapping; nlGC: nonlinear Granger Causality; ML: Machine Learning.

Region	GC	TE	CCM	nlGC	ML
Europe	Zanin (2015) Pastorino and Zanin (2021) Zanin (2021) Belkoura and Zanin (2016)	Zanin (2015)	–	–	Chen et al. (2024)
China	Zanin et al. (2017) Du et al. (2018) Xu and Zhang (2022) Zhang et al. (2022) Pastorino and Zanin (2023)	Xiao et al. (2020) Chen et al. (2023)	Guo et al. (2022)	Jia et al. (2022)	Chen et al. (2024) Li et al. (2023) Kang et al. (2023) Tan et al. (2022)
United States	Mazzarisi et al. (2020) Bombelli and Sallan (2023)	Chen et al. (2023)	Guo et al. (2022)	–	–

3. Data and methods

3.1. Data sources

In this work we considered data extracted from three different sources describing historical air transport operations, each one of different spatial and temporal characteristics. In this study, we restrict the analysis to the years 2015–2018.

The first one is EUROCONTROL's R&D Data Archive, a public repository of historical flights made available for research purposes and freely accessible at https://www.eurocontrol.int/dashboard/rnd-data-archive. It includes information about all commercial flights operating in and over Europe, as well as their flight plans and radar trajectories. The arrival delay for each flight has then been calculated as the difference between the actual and the scheduled landing time. An inherent limitation of this repository is that it only provides data for four months each year (March, June, September, and December).

Secondly, we have considered the Reporting Carrier On-Time Performance database of the Bureau of Transportation Statistics (BTS), U.S. Department of Transportation, freely accessible at https://www.transtats.bts.gov. It contains information about flights operating in US airports, and directly provides the departure and landing delays as pre-calculated features.

While the previous data sets are provided by official organisations and can therefore be regarded as validated records of delays, they are also spatially limited to specific regions, i.e. Europe and the United States. To expand the scope of the analysis, we also consider data from FlightRadar24 (FlightRadar24, 2025b), an internet-based platform that aggregates flight tracking information from multiple sources - denoted as FR24 in what follows for the sake of conciseness. These include primarily crowdsourced data collected via ADS-B receivers operated by volunteers, as well as satellite-based ADS-B systems. Collectively, these sources enable the tracking of over $200,000$ flights per day worldwide (FlightRadar24, 2025a). Most flights ($\sim 80\%$) in the data include scheduled and actual departure and arrival times, from which we calculate the corresponding delays. Since FR24 is not an official source, its information does not necessarily coincide with that from EUROCONTROL or BTS, reflecting a trade-off between broader coverage and differences in data capture philosophy, detailed in section 3.4.

In this work, the three data sets are used in a two-stage and non-overlapping way. First, EUROCONTROL and BTS are used exclusively to develop and validate the two proposed frameworks (microscale and macroscale) on the well-documented Europe-US pair (Section 4). Second, once the methodology has been validated, the global analyses (Sections 5 and 6) are carried out using the FR24 data set only, including for Europe and North America. At no point are flights from different sources merged or combined for the same analysis, each region is always represented by a single data source in a given stage and FR24 data is never mixed with EUROCONTROL or BTS records.

In order to simplify the representation of the global patterns of delay propagation, we have divided the world into nine regions based on the first letter of airports' ICAO code (International Civil Aviation Organization, 2024); each region is then represented by the 50 busiest airports within it, according to the number of flights. The rationale of focusing on the busiest airports is two-fold. On the one hand, it ensures that the analysis captures the primary nodes that dominate regional and intercontinental traffic flows, as these airports serve as effective proxies of the corresponding regional networks. On the other hand, the analyses described below require sufficient quantities of operations to achieve statistically significant results; including low-traffic airports would mask the delay propagation signal and reduce such significance. In short, the use of 50 airports strikes a balance between including enough information to support the analyses, while reducing the risk of such information being noisy or irrelevant. An analysis of the sensitivity of results to the number of airports is included in Supplementary Material, see Figs. 11 and 12. The resulting regional classification scheme is summarised in Table 2; see also Tabs. 1 to 9 in Supplementary Material for the full list of airports.

3.2. Data pre-processing

The previous data sets of flights have been transformed into time series of average delays, where each value $D_i(d, h)$ represents the average delay of all flights landing at airport i at day d and at hour h. It is important to note that these time series are highly non-stationary, i.e. they have trends and seasonalities that are the result of explicit dependencies on time. To illustrate, delays are usually higher during peak hours, when airports and airspaces are saturated; but fewer flights are delayed in the early hours of the morning, especially due to the lack of reactionary delays (The European Organisation for the Safety of Air Navigation, 2012; Campanelli et al., 2016; Monmousseau et al., 2019; Pamplona and Alves, 2020; Cai et al., 2021). Delays are also more pronounced in the summer months, due to both the higher volume of traffic, but also to seasonal weather events (Abdel-Aty et al., 2007; The European Organisa-

Table 2
Definition of the nine world regions considered in this study. The last column reports the set of global ICAO regions (International Civil Aviation Organization, 2024) comprising them. Note that LL* and LT* are excluded from EU and counted in WA; also, GC* is excluded from AF and counted in EU. See Tabs. 1 to 9 in Supplementary Material for the full list of airports.

Region	Code	ICAO Codes
Africa	AF	D*, F*, G*, H*
Europe	EU	E*, L*, GC*
North America	KC	C*, K*, P*
Central America	MT	M*, T*
Oceania	OC	A*, N*, Y*
South America	SM	S*
South Asia	VW	V*, W*
Western and North Asia	WA	LL*, LT*, O*, U*
Central Asia	ZR	R*, Z*

tion for the Safety of Air Navigation, 2024; Performance Review Commission, 2024). This can create biases when the time series are analysed using the techniques described below, as standard statistical tests (including the t- and F-tests) and models (like the Auto Regression Moving Average, ARMA) may yield wrong results; up to the point that propagation of delays can be detected even when randomly shuffling the time series, i.e. when destroying any temporal structure, if no appropriate detrending procedures are applied (Olivares et al., 2025). While it has to be acknowledged that any data manipulation can inadvertently delete relevant information, this is a price that has to be payed in exchange of minimising the number of false positive results, see Sec. “Impact of detrending” in Supplementary Material for further details.

In order to detrend and normalise the delay time series, several alternatives have been evaluated in the literature, with the one based on Z-Scores being the best performing one (Olivares et al., 2025). Such detrend procedure, which is here applied for each airport, is defined as:

$$D'(d,h) = \frac{D(d,h) - \langle D(\cdot,h)\rangle}{\sigma(D(\cdot,h))}, \tag{1}$$

where $D'(d,h)$ is the normalised delay for day d at hour h, $D(d,h)$ is the observed delay, $\langle D(\cdot,h)\rangle$ is the mean delay across all days of the month at hour h, and $\sigma(D(\cdot,h))$ is the corresponding standard deviation. The normalised value $D'(d,h)$ quantifies how many standard deviations the observed delay diverges from the mean value for the specified hour. In other words, a positive value indicates that the delay is greater than expected (i.e., the system performed worse than usual), while a negative value indicates a better-than-expected performance at that time of the day. Most importantly, when the behaviour is in line with what is expected, $D'(d,h)$ is close to zero, irrespective of the original average delay $D(d,h)$.

The performance of this detrending approach is evaluated in Supplementary Material. Specifically, Figs. 1 and 2 report a comparison of the p-values for different time series, both before and after the detrending, as obtained through two stationarity tests: the ADF and the KPSS. All detrended time series are classified as stationary by both tests in a statistically significant way. We further analysed how the detrending may affect the detection of delay propagation instances, see Figs. 3 and 4 in Supplementary Material, and accompanying text. We specifically illustrate how the use of raw (non-detrended) time series may lead to the detection of causal relationships also in the null model, and hence in the reduction of the statistical significance of results; and how no spurious relationships are created by the detrending process.

3.3. Granger causality

To examine the potential causal relationship between two detrended time series, we employ the Granger Causality (GC) test (Granger, 1969). This method evaluates whether the history of one time series (denoted B) provides statistically significant information for predicting the future values of another time series (A), beyond the information contained in the past values of A alone. In other words, if incorporating the past of B allows us to perform better predictions about A, we say that B ‘Granger-causes’ A. Due to the way causality is here tested, the Granger test is usually classified as a predictive causality one (Diebold, 1998). Consider a universe U that represents all elements relevant to a given system, including both observable and hidden variables. Within U, let A and B be two elements described by time series a and b, respectively; these have to be stationary, as previously discussed, and regularly sampled. B is then said to ‘Granger-cause’ A if:

$$\sigma^2(a|U^-) < \sigma^2(a|U^- \setminus b^-), \tag{2}$$

where $\sigma^2(a|U^-)$ represents the error in forecasting the time series a using the past information from the entire universe U and $\sigma^2(a|U^- \setminus b^-)$ represents the error when the information about the past of b is excluded. Note that usually the presence of other

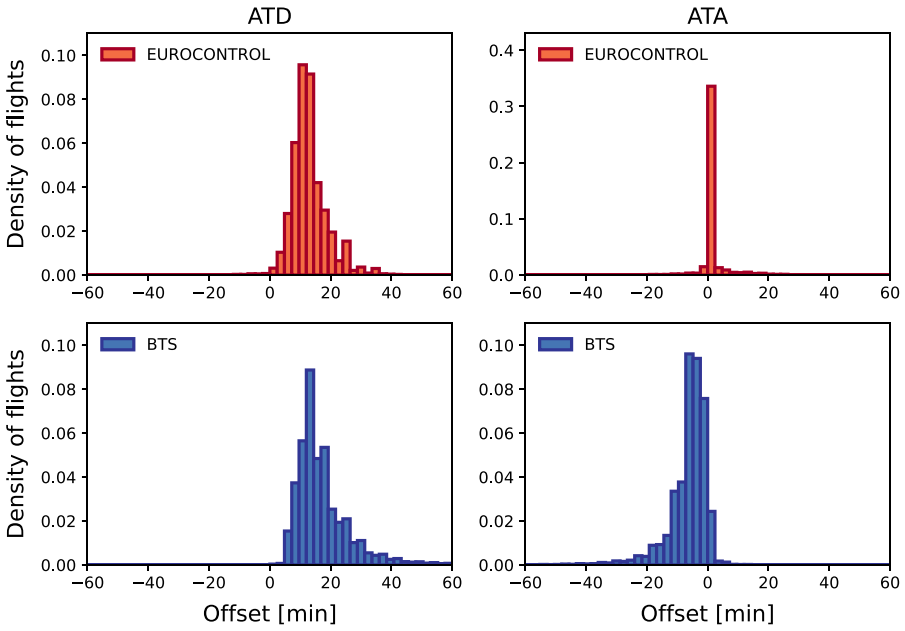

Fig. 1. Time offsets between FR24 and official data sets for matched flights (joined on origin, destination, and STD) in June 2019. Top and bottom panels respectively correspond to the comparison of EUROCONTROL (orange bars) and BTS (blue bars) flights, against FR24. Left and right panels report results for respectively the actual time of departure (ATD) and actual time of arrival (ATA). Data: EUROCONTROL, BTS and FR24. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

elements in the universe is disregarded, to only focus on the pairwise relationship between the two elements under study; the previous condition then simplifies to $\sigma^2(a|b) < \sigma^2(a)$.

The forecast outlined in Eq. (2) can be implemented using various algorithms. One of the most common and straightforward approaches is the autoregressive-moving-average (ARMA) model. Within this framework, the analysis typically involves applying two linear models to the data, known as the restricted and unrestricted regression models, respectively:

$$a_t = C \cdot a^m_{t-l} + \epsilon_t, \tag{3}$$

$$a_t = C' \cdot (a^m_{t-l} \oplus b^m_{t-l}) + \epsilon'_t, \tag{4}$$

where l is the time lag between the time series a and b, m denotes the model order, $\oplus$ represents the concatenation of column vectors, C and C' are matrices containing the model coefficients, and ϵ_t and ϵ'_t are the residuals of the respective models. Here, the restricted model includes only the past values of the series a, while the unrestricted model incorporates both the past values of a and the additional time series b. Then, Eq. (2) is commonly expressed in terms of the residual variances as:

$$\sigma^2(\epsilon'_t) < \sigma^2(\epsilon_t). \tag{5}$$

To evaluate the statistical significance of this improvement and determine whether the inclusion of b has a meaningful impact on the prediction, an F-test is performed. Specifically, the F-test assesses whether the coefficients in C' associated with the series b are significantly different from zero. The test provides a p-value, which ranges between 0 and 1. If the p-value is close to 0, then b 'Granger-causes' a strongly.

The use of the GC test has here been motivated by the fact that it is conceptually and computationally simple: it does not require long time series to achieve statistically significant results, and the interpretation of the causality it defines is quite intuitive. Not surprising, this test has been at the foundation of many of the works described in Section 2. At the same time, it is important to highlight some limitations. On the one hand, it is designed to detect linear relationships - or, to be more exact, the linear part of the relationship (Maziarz, 2015); this implies that some highly nonlinear propagation instances may be underestimated. Secondly, it is well-known that the GC test is sensitive to the presence of confounding effects (Grassmann, 2020) - although this is partially mitigated by the detrending approach of Section 3.2. Finally, as recognised by Granger himself (Granger, 1988), the test does not directly measure causality, as the latter requires intervention.

3.4. Data sets comparison

As described in Section 3.1, EUROCONTROL (EUROCONTROL, 2025a) and BTS (Bureau of Transportation Statistics, U. S. Department of Transportation, 2025) provide official and publicly available data sets describing flight operations in Europe and the

Table 3
Comparison of timestamp definitions across the different data sources. STD = Scheduled Time of Departure, ATD = Actual Time of Departure, STA = Scheduled Time of Arrival and ATA = Actual Time of Arrival. Gate refers to parking brake released/set times and Runway refers to takeoff/landing times.

Source	STD	ATD	STA	ATA
ECTL	Gate	Gate	Runway	Runway
BTS	Gate	Gate	Gate	Gate
FR24	Gate	Runway	Gate	Runway

United States, respectively. These sources can be regarded as validated records of delays, but they are geographically and, in the case of EUROCONTROL, temporally limited. To extend the scope of the analysis, we use data from FlightRadar24 (FlightRadar24, 2025b), a commercial platform that offers worldwide coverage. The official EUROCONTROL and BTS data are first employed to test and validate the proposed delay propagation frameworks (see Section 4). Once this validation step is complete, the large-scale global analyses are conducted exclusively with FR24 data, which provide the required coverage. It is therefore necessary to explicitly analyse the characteristics of data collection across the three sources. Extended methodological detail is provided in Sec. "Data sets comparison" of the Supplementary Material.

In airline and airport operations, several distinct events can serve as temporal references, depending on the aspect of the operation being described. For instance, Off-Blocks refers to the moment when an aircraft releases its parking brake and departs from the stand, whereas On-Blocks denotes its arrival at the stand and setting of the parking brake. These are typically referred to as *gate times*. By contrast, *runway times* are tied to aircraft movements on the runway, e.g. the moment of takeoff for departure and the touch-down for arrival. Different data sets adopt different conventions regarding which of these references is used. EUROCONTROL (EUROCONTROL, 2025a) defines scheduled and actual departures (STD, ATD) as Off-Blocks events, i.e. *gate times*, while scheduled and actual arrivals (STA, ATA) are measured at landing, i.e. *runway times*. BTS (Bureau of Transportation Statistics, U. S. Department of Transportation, 2025) applies a single rule: both departures and arrivals are recorded as *gate times*, with arrival corresponding to the On-Blocks event when the parking brake is set. FR24 (FlightRadar24, 2025a) combines both philosophies: scheduled departure (STD) is reported as a *gate time*, the actual departure (ATD) is reported as a *runway time* (takeoff), the scheduled arrival (STA) is again a *gate time*, and finally the actual arrival (ATA) is a *runway time* (landing). Scheduled times reflect airline planning practices, typically referencing gate operations and incorporating schedule buffers, whereas actual times rely on ADS-B surveillance data that capture runway events with greater precision. A summary of these conventions is provided in Table 3.

A common reference must be identified in order to compare the three data sets. This is provided by the Scheduled Time of Departure (STD), which in all sources is reported as a *gate time*. However, another intrinsic difference lies in the coverage: FR24 consistently reports fewer flights than the official sources, due to coverage limitations. Most of the missing records correspond to regional activity, while intercontinental and interregional flights, departing from or arriving to major hubs, are reliably tracked. For June 2018, by matching flights based on origin, destination, and STD, we linked approximately $\approx 85\%$ of EU flights and $\approx 90\%$ of US flights between the official and FR24 data sets. These matched flights allow for a direct comparison of time stamps. Fig. 1 illustrates the resulting differences. For the Actual Time of Departure (ATD), both FR24-EUROCONTROL and FR24-BTS comparisons yield distributions with a positive mean, consistent with an average taxi-out time of ≈ 10 minutes. The tail extends up to ≈ 40 minutes, a plausible range for large intercontinental hubs subject to congestion. In the US, the tail of the distribution is longer, reflecting higher variability in taxi-out times in the absence of strong gate-holding regulations (ATFM slots). For the Actual Time of Arrival (ATA), the FR24-EUROCONTROL offsets are centred at zero, as both sources adopt the same runway-based definition. In contrast, the FR24-BTS offsets exhibit the reversed effect, with a mean of ≈ -7 minutes, due to BTS using gate times (On-Blocks) and FR24 using runway times (landing). The shorter tail in this case reflects the absence of systematic queuing on arrival. Scheduled Time of Arrival (STA) is not analysed, as differences between sources are dominated by nonlinear factors such as schedule buffers in FR24 and flight plan revisions in the official data.

The previous analysis highlights one important point. While the delay time series obtained from EUROCONTROL, BTS and FR24 data differ in absolute terms, these differences do not stem from inconsistent underlying data, but rather to the particular definitions and conventions of each source. In other words, although individual values differ, they represent the same underlying processes. The relevant question, however, is whether such differences alter the observed propagation behaviour. As shown in Fig. 10 of Supplementary Material, propagation patterns in FR24 and EUROCONTROL official data, while not individually identical, are highly correlated. This result indicates that the specific timestamp definitions shift the baseline of measured delays without distorting the systemic relationships driving their propagation, with FR24 providing the necessary global coverage to extend the analysis beyond the geographical and temporal limits of the official data sets.

4. Two frameworks for assessing long-range delay propagation

Previous GC applications in air transport have evaluated airport-level causality within a single region (Zanin, 2015; Zanin et al., 2017; Du et al., 2018; Pastorino and Zanin, 2021; Jia et al., 2022). We here generalise this idea to the interregional case at two

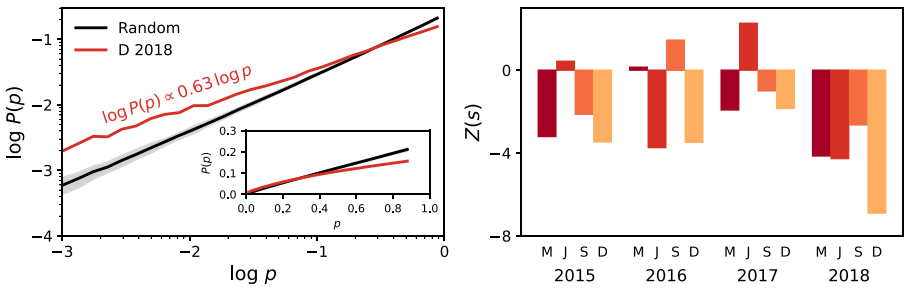

Fig. 2. Evaluating the propagation between two regions. (Left) Probability distribution of the obtained p-values, in the direction EU → US, for the original time series (red solid line), and for 100 random shuffled versions of the same (grey band, with the black line representing the median). (Right) Evolution of the coupling through time, as measured by the Z-Score of the slope $Z(s)$. *M*: March. *J*: June. *S*: September. *D*: December. Data: EUROCONTROL and BTS. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

scales: a microscale approach that pools over many airport pairs to identify weak long-range effects, and a macroscale approach that applies region-to-region GC to detrended regional time series. The microscale framework is sensitive to localised phenomena, identifying sparse but significant causal links by analysing the tail of the pooled p-value distribution. This increases its ability to detect rare but high-impact events. In contrast, the macroscale approach is designed to uncover broad, coherent signals with a lower signal-to-noise ratio. To illustrate these two approaches in a clear and reproducible way, we focus on two regions: Europe and the United States (not included in Table 2, see Tab. 10 in Supplementary Material). These regions were selected because they provide detailed and publicly available regulatory data sets-EUROCONTROL for Europe (EUROCONTROL, 2025a) and BTS for the United States (Bureau of Transportation Statistics, U. S. Department of Transportation, 2025) (see Section 3.1)-which make them well-suited for demonstrating the used methodology. Throughout this section, all results are obtained exclusively from these two official data sets, i.e. FR24 is not used in this methodological validation stage.

4.1. The microscale approach

The microscale approach captures how delays between individual airports are interlinked. We compute the delay time series for each airport and apply a GC test to every ordered pair of airports. A distribution of p-values is obtained across all airport pairs and lags (50 airports x 50 airports x 24 lags). From this distribution, we fit a power law in the range 10^0 to 10^{-3} and extract its exponent, denoted as s_{real}. By design, s_{real} corresponds to the slope of the linear fit on a log-log plot, as shown in Fig. 2 (Left). To assess significance, we repeat the same fitting procedure on the p-value distributions obtained from randomly shuffled versions of the data, resulting in a set of exponents denoted as s_{rnd} and represented by the black line with a grey band in Fig. 2 (Left). The shuffling is performed at the day level: we keep the internal temporal structure of each day intact, but we randomly permute the order of the days in the time series. The difference between the slopes indicates that airport dynamics are not independent and that propagation occurs between the two regions, although this effect is too weak to be detected at the level of individual airports, given that the number of intercontinental flights is much smaller than the number of domestic ones.

To illustrate why a standard approach would not suffice in this case, note that the probability of finding a pair of airports between EU and US with a given p-value, for flights operated in December 2018 fixing a threshold α of 10^{-2}, would yield 3.4% (2,018 pairs) with that or smaller p-value. A correction for multiple comparisons would further be required; as $n = 49 \times 50$ pairs of airports are tested, a Bonferroni correction would yield a corrected $\alpha^* = \alpha/n = 4.082 \times 10^{-6}$, where only 0.01% (9 pairs of airports) fulfil such condition. In short, the small amplitude of the delay propagation signal we aim at detecting precludes the use of a direct analysis of pairs of airports.

Under the null hypothesis of no propagation, the two slopes s_{real} and s_{rnd} should be similar. This is quantified by calculating the Z-Score between both, i.e. $Z_s = [s_{real} - \mu(s_{rnd})]/\sigma(s_{rnd})$, with μ and σ respectively representing the mean and standard deviation, see Fig. 2 (Right). A negative Z-score indicates an excess of small p-values in the real data, consistent with weak propagation.

We test this microscale slopes method on a synthetic model. We specifically consider two coupled dynamical systems, generating time series whose values are defined as:

$$x_a(t) = \mathcal{N}(0,1), \tag{6}$$

$$x_b(t) = (1-\gamma)\mathcal{N}(0,1) + \gamma x_a(t-1). \tag{7}$$

x_a is thus a sequence of normally-distributed random numbers, while x_b is driven, through a linear coupling with strength γ, by the past of x_a. When sufficiently long time series are generated, here of 100 points, a GC test should be able to detect the $x_a \to x_b$ causality whenever the coupling γ is strong enough. This is confirmed by the left panel of Fig. 3, reporting the fraction of tests with a p-value < 0.01 as a function of γ. Reliable detection occurs only for $\gamma > 0.3$. As an alternative, we generate 10^3 pairs of time series with the same dynamics, for then evaluating the exponent k of a power-law fit of the obtained p-values, and finally its Z-Score, as previously described. The resulting Z-Scores, in the right panel of Fig. 3, confirm that this alternative method is able to detect the

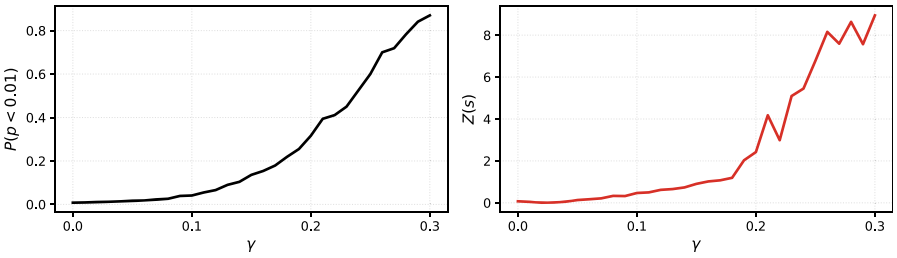

Fig. 3. Slope as a proxy of coupling strength. (Left) Evolution of the fraction of pairs of time series yielding a statistically significant p-value as a function of the coupling γ in the synthetic model. (Right) Associated slope Z-score $Z(s)$ in the p-value distribution fit as a function of γ. Data: Synthetic. .

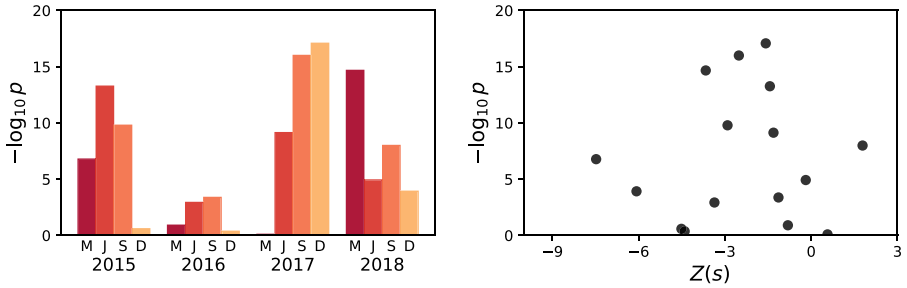

Fig. 4. Comparison between the macroscale and microscale approaches in detecting delay propagation. (Left) Evolution of the p-value p obtained by the macroscale approach for the EU→US propagation, analogous to the right panel of Fig. 2. (Right) Scatter plot of the p-value p from the macroscale approach against the Z-Score of the slope $Z(s)$ from the microscale approach, across all months included in the left panel. Data: EUROCONTROL and BTS.

presence of a causal relationship. The Z-score is zero for $\gamma = 0$, confirming the absence of false positives. The method is more sensitive, a Z-score of 2 (approximately corresponding to $p = .01$) is obtained for $\gamma \approx 0.2$, whereas the fraction of significant tests remains only slightly above 0.2. This method provides a single summary metric for a set of airports, simplifying the interpretation of propagation between regions.

4.2. The macroscale approach

We now introduce the macroscale approach, which focuses on the aggregate dynamics of delays within regions, and compare it to the microscale framework. This is done by firstly creating new time series representing the delay evolution of the entire region. In the microscale case, individual time series $D_i(d, h)$ were generated for each airport i inside a region. On the other hand, here a single time series $D_r(d, h)$ is required for each region r, representing the average delay of all flights landing at any airport inside that region at day d and at hour h; the result is then detrended using the same procedure described in Section 3.2. Note that each value within these time series D_r is created by averaging the delay of all flights landing within a given hour, irrespective of the landing airport. The final average delay thus takes into account the relative weight of each airport, as airports with more operations will contribute more to the average. After detrending, these series represent the evolution of delays at the regional scale rather than simple averages of airport series.

Next, the causality between a pair of regions is calculated by applying the GC test on the corresponding time series, the result being a p-value. In order to obtain results that are comparable with those of the microscale approach, the GC test has been applied to the data corresponding to individual months; when results for longer time windows are needed, these are obtained as the median of the p-values observed in the considered period. The median, denoted $\tilde{p}$, is not a p-value but serves as a proxy for the strength of connectivity between two regions. In other words, the smaller $\tilde{p}$, the stronger delays are expected to propagate.

As the name implies, propagation patterns detected with the macroscale approach describe global dynamics spanning across all airports in a region, as opposed to localised responses to changes in delays. The two approaches may provide substantially different results. To illustrate, if delays propagate between two regions by only affecting two airports in them, such propagation will trivially be detected by the microscale approach; but may be undetectable in the macroscale analysis, as those delays are averaged with all other delays in the system. Conversely, suppose a case in which the propagation results in a small increase in the delays observed in all airports: the propagation will be easy to detect at a macroscale level, but may be too weak to be detected when considering pairs of airports.

To illustrate the relationship between the two methodologies, Fig. 4 compares the macroscale and microscale approaches for US→EU propagation. The left panel shows the time evolution of the macroscale p-value, while the right panel directly compares this to the Z-score of the microscale slope in a scatter plot. As anticipated, the macroscale p-value, $\tilde{p}$, and the microscale Z-score, $Z(\tilde{s})$,

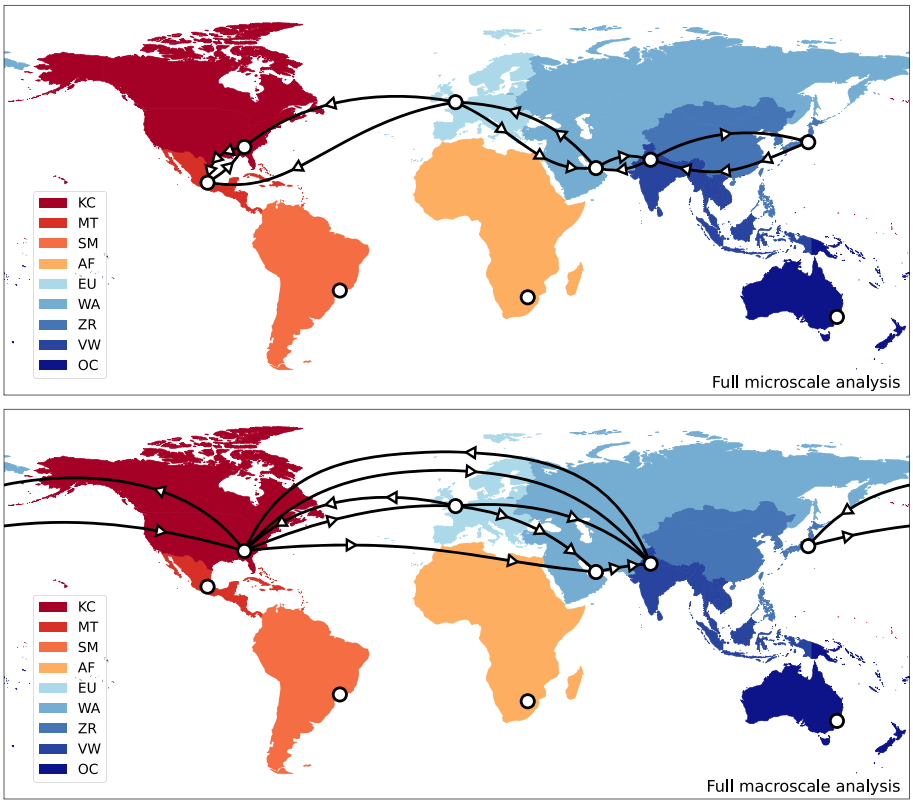

Fig. 5. Global delay propagation maps displaying the top 10 causal links between regions, when calculated using the whole data set. Top and bottom maps respectively correspond to the microscale and macroscale analyses. See Figs. 1 and 2 in Supplementary Material for equivalent plots for the summer and winter seasons. Data: FR24.

exhibit no significant correlation. This is quantitatively confirmed by a $R^2 \approx 0.004$, indicating that the two estimators are largely orthogonal and thus sensitive to complementary features within the data. Given this independence, subsequent analyses use both methodologies in parallel.

5. Global delay propagation patterns

The two proposed methodologies enable the analysis of delay propagation at the global scale, beyond individual airport pairs. We use FR24 data from 2015 to 2018, partitioned into the regions defined in Table 2. In this and the following section, all results are based solely on FR24; EUROCONTROL and BTS are not used in the global propagation analysis. Pairs of regions are examined to identify significant propagation, displayed in the global propagation maps of Fig. 5, with the top and bottom panels corresponding to the micro- and macroscale approaches, respectively. We there depict the 10 strongest directional links between pairs of regions, i.e. for the largest $Z(\tilde{s})$ (microscale, top) and the smallest $\tilde{p}$ (macroscale, bottom). The number of arrows in each link is proportional to the strength of causality - i.e. proportional to $Z(\tilde{s})$ and inversely proportional to $\tilde{p}$. Equivalent plots for summer and winter seasons are provided in Figs. 1 and 2 of the Supplementary Material.

According to the microscale method, a strong bidirectional coupling is observed between KC (North America) and MT (Central America), with the KC→MT direction exhibiting the most pronounced long-range causality. Additionally, the EU (Europe) region propagates delays toward both KC and MT. A prominent intercontinental bridge emerges linking EU, WA, VW and ZR (Europe, Western Asia, South Asia and Central Asia), likely reflecting the crucial role of Gulf-region airports in Europe-Asia air traffic dynamics. The connections within this bridge are largely asymmetric, indicating that delay propagation is stronger in one direction than the reverse. This asymmetry will be examined in further detail in Section 6.

Moving to the macroscale approach, EU and KC emerge as the most connected nodes. EU exhibits the strongest outward delay propagation, primarily affecting KC and WA, while receiving only minor delay propagation from KC. In contrast, KC both transmits and receives weak delay propagation across multiple regions. The link KC→WA is likely associated with the high traffic volume at major Gulf-region airports. As opposed to the microscale approach, the macroscale one does not reveal delay propagation between VW and ZR, effectively breaking the intercontinental bridge.

The differences between the two approaches arise from their distinct underlying mechanisms for detecting causality. As previously discussed (see Section 4.2 and Fig. 4), the microscale approach identifies causality when the delay propagation between two airports

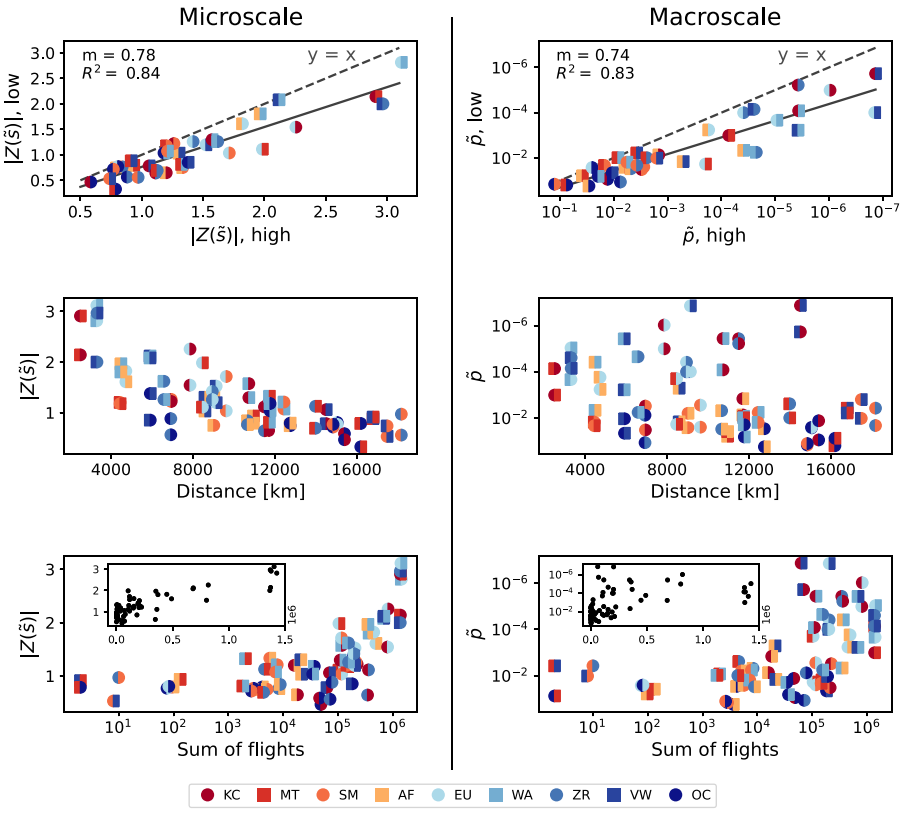

Fig. 6. Scatter plots of the propagation intensity as a function of different metrics. Left and right panels respectively correspond to the microscale and macroscale methods. (Top) Comparison of higher versus lower median slope Z-scores ($Z(\tilde{s})$ and median p-values ($\tilde{p}$), for each pair of regions. (Middle) Absolute value of the metrics against the distance between the region's centroids. (Bottom) Absolute value of the metrics against the accumulated number of flights, from 2015 to 2018. The insets in the latter case report the same information in a linear scale. The left and right halves of each dot represents the origin and destination regions, see colour and shapes in the legend. Data: FR24..

exceeds the background noise level. In contrast, the macroscale approach considers a region as a whole, with delays (both internal and propagated) being evaluated across all airports at the same time. The microscale method is most sensitive when a few airports strongly propagate delays, whereas the macroscale method is favoured when many airports propagate delays weakly. This difference can be visualised, for instance, in the case of regions VW and ZR; as the propagation of delays is centred about a handful of airports, the macroscale approach does not detect a strong-enough signal.

The complete set of links is further analysed in Fig. 6, with left (respectively, right) panels corresponding to the microscale (macroscale) approach. The top panels show scatter plots of the smallest value of the median slope Z-score $Z(\tilde{s})$ (left), and of the $\tilde{p}$ (right), as a function of the largest one. Each point represents a region pair (colours and shapes are given in the legend) and indicates the strongest direction of propagation. Since the slope differs from unity in both cases, the propagation is asymmetric. In both cases, a strong linear correlation is observed, lying close to the main diagonal. This suggests that the propagation of delays in the global aviation network is mostly bi-directional, albeit with varying magnitudes depending on the direction. The observed asymmetry is particularly pronounced in regions with strong causal connections.

The middle panels show the relationship between causality strength and the distance between region centroids, while the bottom panels display the dependence on the number of flights. In both cases and for the microscale approach, a quasi-linear trend emerges, in which the causality is stronger for pairs of regions near in space and connected by many flights. These correlations are not observed in the macroscale approach, as illustrated in Fig. 5. Since the microscale approach is directly influenced by individual airports, it is possible that the analysis is inherently constrained by the physical distance of flights to and from each airport, as well as the number of flights. Thus, only airport pairs connected by multiple flights contribute significantly to the metric. In contrast, the macroscale approach aggregates data from all airports, incorporating long-haul flight delays into a single time series, thereby amplifying their influence.

6. The role and importance of individual airports

We now examine the role of individual airports in the delay propagation process through an approach analogous to wrapper feature selection methods in machine learning, in particular Recursive Feature Elimination (RFE), that identifies the most relevant

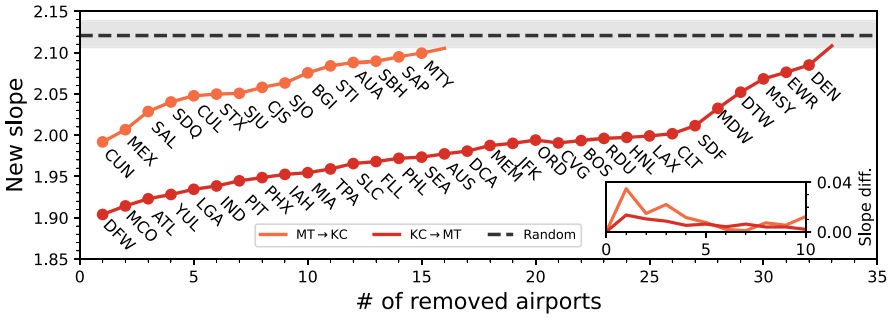

Fig. 7. Evolution of the slope s as airports are iteratively removed in descending order of their impact on the detected delay propagation. The solid lines indicate the slope at each step, in orange for MT→KC and in red for KC→MT; the dashed line represents the slope obtained from randomly permuted data. The inset shows the evolution of the difference in slope between consecutive steps. Data: FR24. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 4
High-impact airports across all region pairs. Percentages report the change in slope when the corresponding airport is eliminated from the analysis. See also Tab. 12 in Supplementary Materials for full results. Data: FR24.

Region Pair	Full	Summer	Winter
KC → MT	DFW, Dallas: 5.7% MCO, Orlando: 11.0% ATL, Atlanta: 15.3%	DFW, Dallas: 6.0% MCO, Orlando: 10.6% ATL, Atlanta: 14.4%	MCO, Orlando: 6.2% TPA, Tampa: 11.6% FLL, Fort Lauderdale: 17.4%
MT → KC	CUN, Cancun: 8.8% MEX, Mexico: 25.9% SJU, San Juan: 35.0%	CUN, Cancun: 21.3% MEX, Mexico: 30.3% SAL, San Salvador: 43.8%	GDL, Guadalajara: 19.3% VER, Veracruz: 25.8% MBJ, Montego Bay: 32.1%
EU → KC	BUD, Budapest: 5.8% AMS, Amsterdam: 8.7% BCN, Barcelona: 11.9%	LIN, Milan: 13.7% MRS, Marseille: 17.2% BRU, Brussels: 19.7%	MAD, Madrid: 7.4% BHX, Birmingham: 15.0% DUB, Dublin: 22.0%
KC → EU	STL, St. Louis: 16.2% ORD, Chicago: 20.6% PHX, Phoenix: 20.7%	BWI, Washington: 14.6% YVR, Vancouver: 22.8% LGA, New York: 27.5%	FLL, Fort Lauderdale: 13.4% IND, Indianapolis: 23.7% RDU, Raleigh: 34.8%
VW → ZR	MFM, Macau: 17.4% SIN, Singapore: 23.3% HKG, Hong Kong: 31.4%	HKG, Hong Kong: 11.6% BKK, Bangkok: 16.1% PEN, George Town: 20.9%	KUL, Kuala Lumpur: 12.2% HKG, Hong Kong: 28.3% MFM, Macau: 36.4%
ZR → VW	CAN, Guangzhou: 7.1% CTU, Chengdu: 12.7% ZUH, Zhuhai: 14.5%	CAN, Guangzhou: 8.4% SZX, Shenzhen: 13.9% HGH, Hangzhou: 16.8%	NKG, Nanking: 4.1% XMN, Xiamen: 13.4% CAN, Guangzhou: 17.7%

features in a classification task by measuring the decrease in performance when a feature is removed (Guyon et al., 2002). In the case at hand, the importance of one airport is estimated by quantifying how much the propagation is changed (here, how much the Z-Score of the slope is altered) when the corresponding data are not taken into account in the analysis. The rationale is straightforward: if an airport plays a dominant role in the propagation, its removal should significantly reduce the observed propagation signal. For a given pair of regions, the analysis proceeds via a greedy elimination procedure, in which the change in the Z-Score is calculated as the data of each airport are eliminated from the slope calculation; the airport for which the largest change is observed is interpreted as the most important one for the propagation. This process is then iterated, each time excluding the airport detected in the previous step, until a complete ranking is obtained.

The full process is graphically illustrated for the pair MT-KC during summer in Fig. 7; the two lines represent how the slope is modified by the removal of the most influential airport in each step. As expected, each iteration moves the slope towards the value obtained for randomly permuted data, i.e. where no propagation is present. The iterative deletions generate a ranking of airports, ordered from higher to lower importance. The change in slope after each removal quantifies the relative contribution of each airport to the propagation dynamics. The most important airports for the KC→MT are Dallas DFW, Orlando MCO and Atlanta ATL. These airports are hubs of airlines with extensive connections to Mexico and Caribbean countries (see Fig. 8). The opposite direction seems to be driven by Cancún CUN, Mexico MEX, and San Salvador SAL, airports representing common touristic destinations during the summer and/or the capitals of the countries of the MT region. An important difference between both directions of propagation can be found in the magnitude of the change in the slope when the three most important airports are deleted; specifically, for KC→MT the slope changes 15%, compared to 44% in the MT→KC case. Thus, in this direction, many airports contribute to propagation, and removing the three most important ones does not substantially alter the dynamics.

Table 4 further reports the three most important airports in the propagation process between the pairs of regions KC-MT, EU-KC, and VW-ZR, further focusing on the summer and winter seasons. In the cases of KC→MT and MT→KC, results for the whole data set align with those corresponding to summer, suggesting that the main delay propagation originates from the saturation of tourism-related flights. This is further supported by the winter dynamics, where the key airports for KC→MT are in Florida, while those for MT→KC are primarily Mexican national airports. This suggests a shift from tourism-driven travel to a higher prevalence of flights associated with personal or community-related mobility, potentially reflecting demographic and sociological patterns in the regions.

For the EU-KC pair, the key result is the relative small impact of airport removal, indicating that no clear set of airports can reliably be pointed at to explain the delay propagation between these regions. However, an interesting pattern emerges: secondary airports

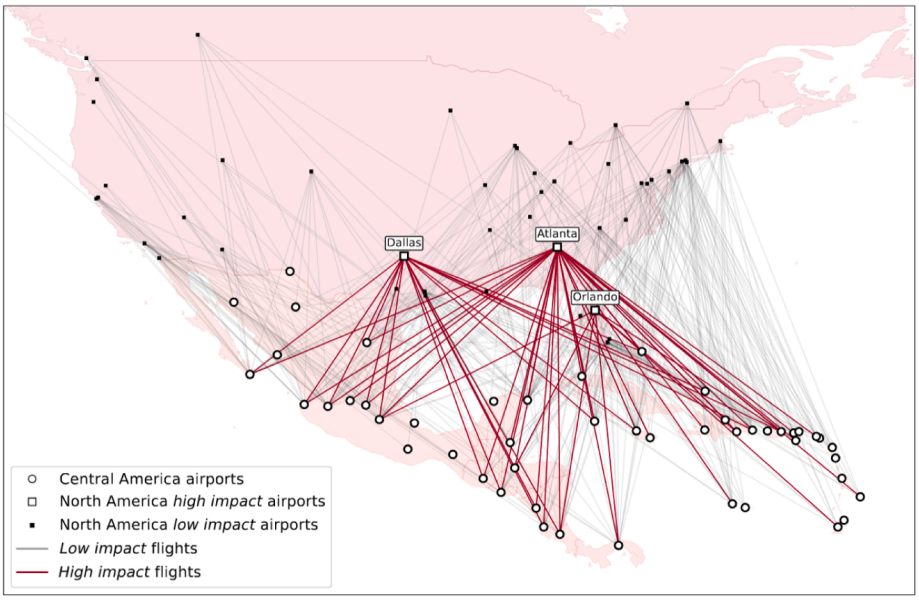

Fig. 8. Depiction of the location of the three most influential airports for the delay propagation for the MT→KC direction. Red links represent direct connections to the KC region departing therefrom; grey links all other connections between the two regions. Data: FR24 for July 2017.

Table 5
Spearman correlation coefficients between airport importance (as measured by the RFE-based slope variation) and traffic volume for all region pairs. Data: FR24..

Reg. Pair	Spear.	Reg. Pair	Spear.	Reg. Pair	Spear.
AF-EU	0.21	EU-VW	−0.01	MT-WA	0.29
AF-KC	0.18	EU-WA	0.29	MT-ZR	0.32
AF-MT	−0.09	EU-ZR	0.18	OC-SM	0.30
AF-OC	0.09	KC-MT	0.50	OC-VW	0.02
AF-SM	0.17	KC-OC	0.20	OC-WA	0.11
AF-VW	0.09	KC-SM	0.17	OC-ZR	0.12
AF-WA	0.20	KC-VW	0.34	SM-VW	0.11
AF-ZR	0.10	KC-WA	0.01	SM-WA	−0.06
EU-KC	−0.21	KC-ZR	0.09	SM-ZR	0.30
EU-MT	0.39	MT-OC	0.11	VW-WA	0.36
EU-OC	0.06	MT-SM	0.12	VW-ZR	0.49
EU-SM	0.27	MT-VW	0.10	WA-ZR	0.30

appear to play a more significant role within these more homogeneous networks. A notable example is LaGuardia Airport in New York, which provides a more representative signal of the behaviour of the US East Coast than JFK itself. Finally, for the pair VW→ZR, delay propagation is primarily driven by airports near the Chinese border. Notably, airports such as Hong Kong and Macao, which belong to China's Special Administrative Regions but start with V ICAO codes, play a significant role. This is an expected outcome given the high volume of flights they operate towards mainland China, as illustrated in Fig. 6. In the opposite direction, ZR→VW, Chinese airports again dominate the propagation of delays outward, which is likewise anticipated due to both the high flight volume and the geographical proximity of the VW region.

From an operational perspective, identifying high-impact airports provides a data-driven foundation for targeted interventions. While the framework pinpoints where delay propagation originates, it does not explicitly attribute these patterns to underlying operational mechanisms such as crew, aircraft, or passenger connection dependencies. Nevertheless, this information remains directly actionable. For instance, the identification of Cancun (CUN) as a major summer propagation source toward North America (Table 4) could enable airlines to refine scheduling practices, such as optimizing turnaround buffers or prioritizing real-time monitoring of operations from this location.

An important question is whether the airport rankings obtained through the RFE procedure simply reflect overall traffic volumes, as, in principle, one might expect airports with higher traffic to dominate the propagation process. To test this, we compared our RFE-based importance metric with three alternative measures of airport centrality: the total number of flights; a centrality metric based on inter-regional traffic flows (Custom Flow Betweenness, CFB); and a centrality metric calculated directly from the Granger-causal propagation networks (Causality-based Betweenness, CaB). The corresponding correlations are shown in Table 5 and illustrated in

Fig. 7 to Fig. 9 of the Supplementary Material. Whereas the comparison with flight counts (Fig. 7 of Supplementary Material) yields moderate positive associations, with maxima of about 0.5 for KC-MT and 0.49 for VW-ZR, both CFB (Fig. 8 of Supplementary Material) and CaB (Fig. 9 of Supplementary Material) show near-zero or weakly negative correlations, with $|\rho| < 0.2$ and non-significant p-values across all region pairs. These results indicate that the RFE-based ranking is largely independent from both static network connectivity, and cannot be explained by the strength of individual causal interaction. In turn, this suggests that airport importance for delay propagation is not trivially determined by traffic volume or by traditional centrality metrics, but rather reflects a distinct, dynamically grounded aspect of network influence.

7. Discussion and conclusions

While the study of the propagation of delays within a region or continent has been central in air transport literature, less attention has been devoted to longer-range instances of the same. One may expect that delays propagate across continents. For example, a late departure from the United States may result in a delayed arrival in Europe, which in turn can force regional flights to wait for connecting passengers and subsequently depart late. Yet, the detection of such propagation is hindered by several factors: intercontinental flights are only a fraction of all flights operating in a given region, and their contribution may be minimal; and these flights are usually long, and have thus more opportunity to recover delays while en route. We address this problem by detecting causality in the evolution of delays in different world regions in a twofold way, comprising both micro- and macroscale analyses. These two approaches are designed to respectively detect instances of propagation that only affected a few airports, and instances involving delays distributed throughout a region. As a consequence, they provide complementary views on the problem, and their results are not always in agreement (see Fig. 5). As an additional challenge, we have observed that different data sets may report delays in a different way (see Section 3.4); while this does not majorly affect our results, it is an issue the practitioner has to be aware of.

Synthesising the results here presented, the first major one is the numerical demonstration that delays indeed propagate through long-range connections - to the best of our knowledge, this was hitherto never been proven. We have observed important propagation instances between many, but not all, regions. As is to be expected, these have strong seasonal variations, reflecting the changes in connectivity during summer/winter. The patterns are largely symmetrical: if region A propagates delays to B, then B also propagates delays back to A, although with varying intensity (see top panels of Fig. 6). Additionally, two regions will suffer a stronger propagation whenever they are close in space, and whenever they are connected by a large number of flights (see respectively middle and bottom panels of Fig. 6). These general trends are common to both the micro- and the macroscale approaches; the only exception is that the relationship with the distance is lost in the latter case. This may indicate that global propagation patterns, i.e. those involving more than a few airports in each region, may be driven by non-local events, such as weather patterns extending across multiple continents through teleconnections (Liu and Alexander, 2007). This hypothesis is further confirmed by the fact that most propagation instances seem to occur at similar latitudes, i.e. may be driven by atmospheric circulation patterns. We were additionally able to detect the key airports in the propagation of delays between pairs of regions, through a process similar to Recursive Feature Elimination; and to highlight how their identity is associated to known traffic patterns (see Table 4).

Regarding the direct operational implications of the methodology and results here presented, it is essential to distinguish between the descriptive capability of the proposed frameworks and prescriptive operational measures. The methods empirically quantify where and between which regions delays propagate, and identify key airport transmission channels (see Table 4). Yet, they do not resolve the underlying causal mechanisms; in other words, they do not allow to identify which airline policy is mostly responsible for such propagations. Despite not revealing the causes of propagation, the approach provides, for the first time, a global and data-driven perspective on how delays spread across continents, making it a valuable instrument for strategic monitoring, risk assessment and network resilience. Airlines and air traffic managers can use these insights to extend disruption surveillance beyond regional networks, and to target interventions at the airports identified as main transmission channels. For example, a European carrier could use seasonal rankings to anticipate delay risks originating from specific African airports; or a US carrier could adjust turnaround or block times for flights departing from high-risk airports, such as CUN in summer or major Florida airports in winter (Table 4). Beyond operational applications, the findings also support policy makers by revealing seasonally and regionally relevant airports from a propagation viewpoint, thereby guiding prioritisation programmes in saturated airspace and airports (Janic, 2009). Notwithstanding these potential operational implications, the validation of any intervention would require additionally analyses, as e.g. micro-scale approaches to evaluate what-if scenarios through large-scale simulations. The proposed methodology can thus be seen as a first necessary step towards a more detailed understanding of the large-scale propagation of delays.

As an empirical study of long-range delay propagation, this work also highlights several challenges, which may be the foundations of future research efforts. Firstly, as already noted, there is a heterogeneity in the way delays are reported in each region, making such studies more challenging. Secondly, while the Granger Causality is a well-established methodology, it is not free from drawbacks - most notably, the lack of sensitivity for nonlinear causality relationships. Many alternatives are available, including nonlinear versions of the GC test (Marinazzo et al., 2008; Zanin, 2024), Transfer Entropy (Staniek and Lehnertz, 2008; Zanin, 2015; Xiao et al., 2020; Chen et al., 2023) or the Convergent Cross Mapping (CCM) (Ye et al., 2015; Tsonis et al., 2017; Guo et al., 2022). It has nevertheless to be noted that these usually require longer time series, which may not be available, and which would hinder the analysis of fast dynamics. In any case, the temporal scales inherent to this type of analysis constitute a significant limitation, as they do not allow the study of system behaviour under specific conditions. Thirdly, we have here evaluated two complementary approaches to the study of long-range delay propagations, respectively focusing on the micro- and macroscale dynamics; yet, a more coherent picture may be obtained by integrating the two approaches into a single method that incorporates information at multiple spatial scales, for instance by analysing the aggregated dynamics of mesoscale groups of airports, similar to a clustering approach.

CRediT authorship contribution statement

Josu Blanco: Writing – review & editing, Writing – original draft, Software, Methodology, Formal analysis; **Antònia Tugores:** Writing – review & editing, Writing – original draft, Software, Formal analysis; **José J. Ramasco:** Writing – review & editing, Writing – original draft, Methodology, Conceptualization; **Massimiliano Zanin:** Writing – review & editing, Writing – original draft, Methodology, Conceptualization.

Data availability

All data sources are open, and links are provided within the manuscript.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Massimiliano Zanin reports financial support was provided by European Research Council. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

J.B. acknowledges economic support by the Conselleria d'Educació i Universitats of the Government of the Balearic Islands (FPI 042/2024). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 851255). This work was partially supported by the María de Maeztu project CEX2021-001164-M funded by the MICIU/AEI/10.13039/501100011033 and also by the APASOS project from MICIU/AEI/10.13039/501100011033 and Fondo Europeo de Desarrollo Regional (FEDER, UE). The authors thank Flightradar24 for granting permission to use their data. Flightradar24 cannot be held liable for the accuracy of the information or for ensuring that the information is up to date at all times. It is thus not liable for any consequences that may arise as a result of possible inaccuracies in the information provided.

Supplementary material

Supplementary material associated with this article can be found in the online version at 10.1016/j.tre.2026.104688.

References

Abdel-Aty, M., Lee, C., Bai, Y., Li, X., Michalak, M., 2007. Detecting periodic patterns of arrival delay. J. Air Transp. Manage. 13 (6), 355–361.
Acharya, K., Olivares, F., Zanin, M., 2024. How representative are air transport functional complex networks? a quantitative validation. Chaos An Interdiscip. J. Nonlinear Sci. 34 (4).
AhmadBeygi, S., Cohn, A., Guan, Y., Belobaba, P., 2008. Analysis of the potential for delay propagation in passenger airline networks. J. Air Transp. Manage. 14 (5), 221–236.
Ajayi, R.A., Friedman, J., Mehdian, S.M., 1998. On the relationship between stock returns and exchange rates: tests of granger causality. Global Finance J. 9 (2), 241–251.
Albert, R., Jeong, H., Barabási, A.-L., 1999. Diameter of the world-wide web. Nature 401 (6749), 130–131.
Albert, R., Jeong, H., Barabási, A.-L., 2000. Error and attack tolerance of complex networks. Nature 406 (6794), 378–382.
Barabási, A.-L., Bonabeau, E., 2003. Scale-free networks. Sci. Am. 288 (5), 60–69.
Barthélemy, M., 2011. Spatial networks. Phys. Rep. 499 (1–3), 1–101.
Baspinar, B., Koyuncu, E., 2016. A data-driven air transportation delay propagation model using epidemic process models. Int. J. Aerosp. Eng. 2016 (1), 4836260.
Beatty, R., Hsu, R., Berry, L., Rome, J., 1999. Preliminary evaluation of flight delay propagation through an airline schedule. Air Traffic Control Quart. 7 (4), 259–270.
Belkoura, S., Peña, J.M., Zanin, M., 2017. Beyond linear delay multipliers in air transport. J. Adv. Transp. 2017 (1), 8139215.
Belkoura, S., Zanin, M., 2016. Phase changes in delay propagation networks. arXiv:1611.00639.
Bombelli, A., Sallan, J.M., 2023. Analysis of the effect of extreme weather on the US domestic air network. a delay and cancellation propagation network approach. J. Transp. Geogr. 107, 103541.
Boswell, S.B., Evans, J.E., 1997. Analysis of Downstream Imports of Air Traffic Delay. Technical Report. Massachusetts Institute of Technology. Lincoln Laboratory.
Brueckner, J.K., Czerny, A.I., Gaggero, A.A., 2021. Airline mitigation of propagated delays via schedule buffers: theory and empirics. Transp. Res. Part E Logist. Transp. Rev. 150, 102333.
Bryan, D.L., O'kelly, M.E., 1999. Hub-and-spoke networks in air transportation: an analytical review. J. Reg. Sci. 39 (2), 275–295.
Bureau of Transportation Statistics, U. S. Department of Transportation, 2025. Reporting carrier on-time performance. https://www.transtats.bts.gov.
Cai, Q., Alam, S., Duong, V.N., 2021. A spatial–temporal network perspective for the propagation dynamics of air traffic delays. Engineering 7 (4), 452–464.
Campanelli, B., Fleurquin, P., Arranz, A., Etxebarria, I., Ciruelos, C., Eguíluz, V.M., Ramasco, J.J., 2016. Comparing the modeling of delay propagation in the US and european air traffic networks. J. Air Transp. Manage. 56, 12–18.
Carlier, S., De Lépinay, I., Hustache, J.-C., Jelinek, F., 2007. Environmental impact of air traffic flow management delays. In: 7th USA/Europe Air Traffic Management Research and Development Seminar (ATM2007). Vol. 2. Citeseer, p. 16.
Chen, G., Fricke, H., Okhrin, O., Rosenow, J., 2024. Flight delay propagation inference in air transport networks using the multilayer perceptron. J. Air Transp. Manage. 114, 102510.
Chen, J., Li, M., 2019. Chained predictions of flight delay using machine learning. In: AIAA Scitech 2019 Forum, p. 1661.
Chen, S., Du, W., Liu, R., Cao, X., 2023. Finding spatial and temporal features of delay propagation via multi-layer networks. Physica A 614, 128526.
Cook, A., Blom, H. A.P., Lillo, F., Mantegna, R.N., Micciché, S., Rivas, D., Vázquez, R., Zanin, M., 2015. Applying complexity science to air traffic management. J. Air Transp. Manage. 42, 149–158.

Cook, A.J., Tanner, G., 2011. European Airline Delay Cost Reference Values. Technical Report. University of Westminster.
Diebold, F.X., 1998. Elements of Forecasting. South-Western College Pub. Cincinnati, OH, USA.
Ding, M., Chen, Y., Bressler, S.L., 2006. Granger causality: basic theory and application to neuroscience. Handbook of time series analysis: Recent Theoretical Developments and Applications, 437–460.
Du, W.-B., Zhang, M.-Y., Zhang, Y., Cao, X.-B., Zhang, J., 2018. Delay causality network in air transport systems. Transp. Res. Part E: Logist. Transp. Rev. 118, 466–476.
EUROCONTROL, 2023. Network Operations Report. Technical Report, EUROCONTROL.
EUROCONTROL, 2025a. Aviation data repository for research. https://www.eurocontrol.int/dashboard/aviation-data-research.
EUROCONTROL, 2025b. EUROCONTROL. https://www.eurocontrol.int.
Feng, D., Hao, B., Lai, J., 2024. Tracing delay network in air transportation combining causal propagation and complex network. International Journal of Intelligent Networks 5, 63–76.
Fleurquin, P., Ramasco, J.J., Eguíluz, V.M., 2013. Systemic delay propagation in the US airport network. Sci. Rep. 3 (1), 1159.
FlightRadar24, 2025a. FR24. https://www.flightradar24.com/about.
FlightRadar24, 2025b. FR24 API. https://fr24api.flightradar24.com.
Gopalakrishnan, K., Balakrishnan, H., 2021. Control and optimization of air traffic networks. Ann. Rev. Contr. Robot. Auton. Syst. 4 (1), 397–424.
Granger, C. W.J., 1969. Investigating causal relations by econometric models and cross-spectral methods. Econometrica, 424–438.
Granger, C. W.J., 1988. Causality, cointegration, and control. J. Econ. Dyn. Control 12 (2-3), 551–559.
Grassmann, G., 2020. New considerations on the validity of the wiener-granger causality test. Heliyon 6 (10).
Guimera, R., Mossa, S., Turtschi, A., Amaral, L.A.N., 2005. The worldwide air transportation network: anomalous centrality, community structure, and cities' global roles. Proc. Nat. Acad. Sci. 102 (22), 7794–7799.
Guo, Z., Hao, M., Yu, B., Yao, B., 2022. Detecting delay propagation in regional air transport systems using convergent cross mapping and complex network theory. Transp. Res. Part E Logist. Transp. Rev. 157, 102585.
Guyon, I., Weston, J., Barnhill, S., Vapnik, V., 2002. Gene selection for cancer classification using support vector machines. Mach. Learn. 46, 389–422.
Hansen, M., 2002. Micro-level analysis of airport delay externalities using deterministic queuing models: a case study. J. Air Transp. Manage. 8 (2), 73–87.
Hiemstra, C., Jones, J.D., 1994. Testing for linear and nonlinear granger causality in the stock price-volume relation. J. Finance 49 (5), 1639–1664.
International Civil Aviation Organization, 2024. ICAO Document 7910: location indicators. International Civil Aviation Organization. Montreal, Canada. 193 ed.
Janic, M., 2009. Concept for prioritizing aircraft operations at congested airports. Transp. Res. Rec. 2106 (1), 100–108.
Jia, Z., Cai, X., Hu, Y., Ji, J., Jiao, Z., 2022. Delay propagation network in air transport systems based on refined nonlinear granger causality. Transportmetrica B: Transp. Dyn. 10 (1), 586–598.
Kang, J., Yang, S., Shan, X., Bao, J., Yang, Z., 2023. Exploring delay propagation causality in various airport networks with attention-weighted recurrent graph convolution method. Aerospace 10 (5), 453.
Kim, M., Bae, J., 2021. Modeling the flight departure delay using survival analysis in south korea. J. Air Transp. Manage. 91, 101996.
Latora, V., Marchiori, M., 2003. Economic small-world behavior in weighted networks. Eur. Phys. J. B-Condensed Matter Complex Syst. 32, 249–263.
Li, C., Mao, J., Li, L., Wu, J., Zhang, L., Zhu, J., Pan, Z., 2024. Flight delay propagation modeling: data, methods, and future opportunities. Transp. Res. Part E Logist. Transp. Rev. 185, 103525.
Li, Q., Guan, X., Liu, J., 2023. A CNN-LSTM framework for flight delay prediction. Expert Syst. Appl. 227, 120287.
Li, C., Qi, X., Yang, Y., Zeng, Z., Zhang, L., Mao, J., 2024. FAST-CA: fusion-based adaptive spatial–temporal learning with coupled attention for airport network delay propagation prediction. Inform. Fusion 107, 102326.
Li, S., Xie, D., Zhang, X., Zhang, Z., Bai, W., 2020. Data-driven modeling of systemic air traffic delay propagation: an epidemic model approach. J. Adv. Transp. 2020 (1), 8816615.
Liu, Z., Alexander, M., 2007. Atmospheric bridge, oceanic tunnel, and global climatic teleconnections. Rev. Geophys. 45 (2).
Lu, Q., Dessouky, M., Leachman, R.C., 2004. Modeling train movements through complex rail networks. ACM Trans. Modeling Comput. Simul. (TOMACS) 14 (1), 48–75.
Marinazzo, D., Pellicoro, M., Stramaglia, S., 2008. Kernel method for nonlinear granger causality. Phys. Rev. Lett. 100 (14), 144103.
Maziarz, M., 2015. A review of the granger-causality fallacy. The J. Philosoph. Econ. Reflect. Econ. Soc. Issues 8 (2), 86–105.
Mazzarisi, P., Zaoli, S., Lillo, F., Delgado, L., Gurtner, G., 2020. New centrality and causality metrics assessing air traffic network interactions. J. Air Transp. Manage. 85, 101801.
McGraw, M.C., Barnes, E.A., 2018. Memory matters: a case for granger causality in climate variability studies. J. Clim. 31 (8), 3289–3300.
Monmousseau, P., Delahaye, D., Marzuoli, A., Féron, E., 2019. Predicting and analyzing US air traffic delays using passenger-centric data-sources. In: ATM 2019, 13th USA/Europe Air Traffic Management Research and Development Seminar.
Newman, M. E.J., 2003. The structure and function of complex networks. SIAM Rev. 45 (2), 167–256.
Olivares, F., Marín-Rodríguez, F.J., Acharya, K., Zanin, M., 2025. Evaluating methods for detrending time series using ordinal patterns, with an application to air transport delays. Entropy 27 (3), 230.
Pamplona, D.A., Alves, C. J.P., 2020. An overview of air delay: a case study of the Brazilian scenario. Transp. Res. Interdiscip. Persp. 7, 100189.
Pastorino, L., Zanin, M., 2021. Air delay propagation patterns in Europe from 2015 to 2018: an information processing perspective. J. Phys. Complex. 3 (1), 015001.
Pastorino, L., Zanin, M., 2023. Local and network-wide time scales of delay propagation in air transport: a granger causality approach. Aerospace 10 (1), 36.
Performance Review Commission, 2024. Observer No. 2 - October 2024. Technical Report, EUROCONTROL.
Pyrgiotis, N., Malone, K.M., Odoni, A., 2013. Modelling delay propagation within an airport network. Transp. Res. Part C Emerg. Technol. 27, 60–75.
Ravasz, E., Barabási, A.-L., 2003. Hierarchical organization in complex networks. Phys. Rev. E 67 (2), 026112.
Ruehle, J., Goetsch, B., Koch, B., 2006. Consequences of feeder delays for the success of a380 operations. J. Air Transp. 11 (2).
Seth, A.K., Barrett, A.B., Barnett, L., 2015. Granger causality analysis in neuroscience and neuroimaging. J. Neurosci. 35 (8), 3293–3297.
Silva, F.N., Vega-Oliveros, D.A., Yan, X., Flammini, A., Menczer, F., Radicchi, F., Kravitz, B., Fortunato, S., 2021. Detecting climate teleconnections with granger causality. Geophys. Res. Lett. 48 (18), e2021GL094707.
Song, C., Guo, J., Zhuang, J., 2020. Analyzing passengers' emotions following flight delays-a 2011–2019 case study on SKYTRAX comments. J. Air Transp. Manage. 89, 101903.
Staniek, M., Lehnertz, K., 2008. Symbolic transfer entropy. Phys. Rev. Lett. 100 (15), 158101.
Strogatz, S.H., 2001. Exploring complex networks. Nature 410 (6825), 268–276.
Tan, X., Liu, Y., Liu, D., Zhu, D., Zeng, W., Wang, H., 2022. An attention-based deep convolution network for mining airport delay propagation causality. Appl. Sci. 12 (20), 10433.
The European Organisation for the Safety of Air Navigation, 2012. CODA Digest - Delays to Air Transport in Europe, April 2012. Technical Report, EUROCONTROL.
The European Organisation for the Safety of Air Navigation, 2024. CODA Digest - All-Causes Delays to Air Transport in Europe. Technical Report, EUROCONTROL.
Tsonis, A.A., Deyle, E.R., Ye, H., Sugihara, G., 2017. Convergent cross mapping: theory and an example. Adv. Nonlinear Geosci., 587–600.
Vỳrost, T., Lyócsa, Š., Baumöhl, E., 2015. Granger causality stock market networks: temporal proximity and preferential attachment. Physica A 427, 262–276.
Wandelt, S., Xu, Y., Sun, X., 2023. Measuring node importance in air transportation systems: on the quality of complex network estimations. Reliab. Eng. Syst. Saf. 240, 109596.
Wong, J.-T., Tsai, S.-C., 2012. A survival model for flight delay propagation. J. Air Transp. Manage. 23, 5–11.
Wu, Y., Yang, H., Lin, Y., Liu, H., 2023. Spatiotemporal propagation learning for network-wide flight delay prediction. IEEE Trans. Knowl. Data Eng. 36 (1), 386–400.
Xiao, Y., Zhao, Y., Wu, G., Jing, Y., 2020. Study on delay propagation relations among airports based on transfer entropy. IEEE Access 8, 97103–97113.
Xie, F., Levinson, D., 2007. Measuring the structure of road networks. Geogr. Anal. 39 (3), 336–356.
Xu, G., Zhang, X., 2022. Statistical analysis of resilience in an air transport network. Front. Phys. 10, 969311.

Ye, H., Deyle, E.R., Gilarranz, L.J., Sugihara, G., 2015. Distinguishing time-delayed causal interactions using convergent cross mapping. Sci. Rep. 5 (1), 14750.
Zanin, M., 2015. Can we neglect the multi-layer structure of functional networks? Physica A 430, 184–192.
Zanin, M., 2021. Simplifying functional network representation and interpretation through causality clustering. Sci. Rep. 11 (1), 15378.
Zanin, M., 2024. Augmenting granger causality through continuous ordinal patterns. Commun. Nonlinear Sci. Numer. Simul. 128, 107606.
Zanin, M., Belkoura, S., Zhu, Y., 2017. Network analysis of Chinese air transport delay propagation. Chin. J. Aeronaut. 30 (2), 491–499.
Zanin, M., Lillo, F., 2013. Modelling the air transport with complex networks: a short review. Eur. Phys. J. Special Top. 215 (1), 5–21.
Zhang, X., Zhao, S., Mei, H., 2022. Analysis of airport risk propagation in Chinese air transport network. J. Adv. Transp. 2022 (1), 9958810.
Zheng, Z., Wei, W., Hu, M., 2021. A comparative analysis of delay propagation on departure and arrival flights for a Chinese case study. Aerospace 8 (8), 212.