Second US-Europe Air Transportation Research and Development Symposium (ATRDS2026)

Too Busy to Depeak: Quantifying the Persistence of Peaked Schedules at European Airports

Josu Blanco-Ortiz, José J. Ramasco
Instituto de Física Interdisciplinar y Sistemas Complejos
IFISC (CSIC-UIB)
Campus Universitat de les Illes Balears,
E-07122 Palma de Mallorca, Spain
{josu, jramasco}@ifisc.uib-csic.es

Massimiliano Zanin
Instituto de Física Interdisciplinar y Sistemas Complejos
IFISC (CSIC-UIB)
Edifici Complex de Recerca de les Illes Balears,
Parc Bit, E-07120, Palma de Mallorca, Spain
mzanin@ifisc.uib-csic.es

***Abstract*—The temporal organisation of airline schedules is shaped by the superposition of individual carrier decisions. Banked schedules concentrate arrivals and departures in waves, supporting transfer connectivity by reducing passenger connection times, while also increasing concurrent demand on airport and airspace resources. In this work, we study this temporal concentration primarily from the airport-level perspective, complemented by a preliminary decomposition by leading airlines. We introduce two complementary indicators of schedule peakiness: the Shannon entropy of the hourly traffic profile and the number of busy hours defined through airport-specific thresholds. Using the COVID-19 pandemic as a reference disruption, we compare pre- and post-pandemic schedules across major European airports at comparable traffic levels. Results show substantial airport-level heterogeneity: Rome Fiumicino and Dublin display flatter post-pandemic aggregate schedules, whereas airports such as London Heathrow, Amsterdam Schiphol and Zurich show persistent or reinforced concentration. The airline-level decomposition suggests that these airport-wide changes reflect carrier-specific shifts in schedule organisation. Overall, we find no evidence of a general post-pandemic depeaking of European airport schedules.**

***Keywords*—airline scheduling, airport operations, depeaking, entropy, COVID-19**

I. Introduction

Since the 1980s, many major airlines have adopted hub-and-spoke networks with banked schedules, i.e. concentrated waves of arrivals and departures designed to facilitate passenger connections [1], [2]. In this context, connectivity refers primarily to the ability of passengers to transfer between inbound and outbound flights with short connection times, thereby expanding the set of feasible itineraries and reducing total travel time. The same temporal concentration, however, also increases concurrent demand for airport infrastructure, air traffic control, gates, ground handling, and other local resources.

A less temporally concentrated alternative is usually referred to as depeaking, in favour of a more "rolling hub" (or "continuous scheduling") strategy. Such approach involves trade-offs: it may reduce simultaneous resource demand and congestion, but it can also weaken connection banks, increase passenger transfer times, and affect overall airline revenue and network performance [3]–[7]. Related studies have also considered airport-side implications [8], [9]. Rome Fiumicino (LIRF) has previously been identified as a case of outbound depeaking following a reduction in peak departure throughput in 2017 [10].

We present preliminary results quantifying changes in airport-level schedule peakiness through two complementary metrics: Shannon entropy of hourly traffic profiles, which is higher when operations are more evenly distributed, and the count of hours exceeding a high-utilisation threshold. The unit of analysis in this work is thus the aggregate airport schedule, complemented by a preliminary decomposition by leading airlines. This does not imply that airports unilaterally determine the temporal structure of operations. Airline decisions, passenger demand, fleet rotations, crew constraints, and network connectivity requirements are central determinants of schedule construction [11]–[13]. Nevertheless, the superposition of individual airline schedules is experienced at the airport as an aggregate temporal load profile.

Using COVID-19 as a reference disruption, we compare pre- and post-pandemic peakiness at comparable traffic volumes across European airports. This allows us to distinguish changes in temporal organisation from changes caused only by traffic recovery. Our analysis reveals divergent patterns: some airports, such as LIRF and EIDW, display higher post-pandemic entropy at similar traffic levels, while others, including EGLL, EHAM and LSZH, maintain or reinforce temporally concentrated patterns. Overall, the preliminary results provide limited evidence of a systematic change in airport-level schedule peaking or depeaking across Europe.

II. Quantifying depeaking

Let us take as starting point the set of operations f scheduled to operate in a given airport in a given day, with $N = |f|$ being the total number of operations (both departures and arrivals). The day is divided into 24 hourly slots, so that we define a vector $\vec{n} = (n_0, n_1, \ldots, n_{23})$, with n_h representing the number of operations in hour h and $\sum_h n_h = N$. Normalising by N gives the hourly fraction of daily operations $\vec{p} = (p_0, p_1, \ldots, p_{23})$ with $p_h = n_h/N$.

The uniformity of the latter distribution is quantified through the Shannon entropy [14], [15]:

$$H = -\sum_{h=0}^{23} p_h \log_2 p_h. \tag{1}$$

Large values of H indicate a more homogeneous hourly distribution of operations across the 24-hour period; small values of H indicate stronger temporal concentration in fewer hours. We additionally capture high-load periods by setting an airport-specific threshold τ, defined as the 90$^{\text{th}}$ percentile of all non-zero hourly counts n_h in the pre COVID-19 period for the given airport. The number of busy hours is then computed as:

$$B = \sum_{h=0}^{23} \Theta(n_h - \tau), \tag{2}$$

with Θ being the Heaviside function (zero if the argument is negative, and one otherwise). Both H and B are analysed conditional on the total daily traffic N. This avoids confounding depeaking with traffic growth or recovery. As airports approach capacity, additional operations may naturally be placed in less busy periods, thereby increasing the apparent uniformity of the schedule. Such a mechanism would not, by itself, constitute depeaking in our framework. We therefore compare pre and post COVID-19 schedules only at comparable values of N, and interpret depeaking as a shift in the relationship between temporal organisation and traffic volume.

We apply the two metrics to two representative examples, Rome Fiumicino (LIRF) and London Heathrow (EGLL), see Fig. 1. As will be shown below, these two airports exemplify two different ways of accommodating traffic, specifically when considering the recovering from the large-scale disruption caused by the COVID-19 pandemic; other airports have dynamics that mostly fit between these two.

The left panels show daily scheduled operations for both airports. Both experienced substantial COVID-19 drops (grey band) but have since recovered pre-pandemic traffic levels. This naturally defines two periods of comparable activity: a pre COVID-19 one, ending March 1^{st} 2020; and a post COVID-19 one, starting December 31^{st} 2021, i.e. roughly corresponding to half-way recovery. The central panels display scatter plots of busy hours B (Eq. 2) versus daily traffic count N. Key differences emerge: for EGLL, pre (blue) and post COVID-19 (red) days overlap at high traffic levels, whereas LIRF shows consistently fewer busy hours post-pandemic at comparable traffic. The right panels show entropy H (Eq. 1) versus N, with similar results: highly overlapping distributions for EGLL, but distinctly higher post COVID-19 entropy for LIRF, indicating more homogeneous distribution of operations throughout the day.

Two contrasting airport-level patterns thus emerge. At EGLL, the pre and post COVID-19 distributions largely overlap at comparable traffic levels, suggesting that the temporal concentration of the aggregate schedule remained similar across the two periods. At LIRF, post COVID-19 days exhibit higher entropy and fewer busy hours at comparable values of N, indicating a shift towards a more even distribution of operations.

EGLL	LFPG	EHAM	EDDF	LEMD	LEBL
EDDM	EGKK	LIRF	LFPO	EIDW	LSZH
EKCH	LEPA	LPPT	ENGM	EGCC	EGSS
LOWW	ESSA	EBBR	LIMC	EDDL	LGAV
EDDT	LEMG	EPWA	LSGG	EDDH	LKPR

TABLE I. List of the 30 airports considered in this study, identified by their ICAO code.

III. Extending the analysis to other airports

We next analyse whether similar patterns emerge for other airports, specifically focusing on operations between 2012 and 2024 for the 30 largest European airports - see the list in Tab. I. Due to limitations in the available dataset, the analysis is restricted to Mondays. This provides one consistent weekly observation per airport, while avoiding the additional heterogeneity that would arise from mixing weekdays with potentially different scheduling patterns.

Extending the analysis beyond two airports requires quantitative metrics to capture the observed differences. As shown in the central panels of Fig. 1, the evolution of busy hours B follows two phases: an initial plateau where traffic is insufficient to generate busy hours, and a subsequent approximately linear increase (illustrated with synthetic data in the left panel of Fig. 2). Basic statistics like the average of B would be misleading, as they include the irrelevant plateau region. We therefore identify the smallest traffic count where $B > 0$ and the maximum traffic count, both pre and post COVID-19 (represented in Fig. 2 by dashed and dotted vertical lines, blue and red respectively). Within the overlap region where $B > 0$ for both periods, we fit two linear regressions with a shared slope constraint. The shared slope makes the comparison explicitly conditional on traffic volume: the relevant quantity is the vertical displacement between the pre and post COVID-19 curves for identical N. This assumption is justified by Fig. 3, which shows that the slopes obtained by fitting the pre and post COVID-19 periods cluster around the main diagonal. The fit yields three parameters: slope β and two intercepts α_{pre} and α_{post}. In the synthetic example of Fig. 2, this produces $\beta \approx 0.048$, $\alpha_{\text{pre}} \approx -34.8$, and $\alpha_{\text{post}} \approx -37.2$.

The difference $\delta_B = \alpha_{\text{post}} - \alpha_{\text{pre}}$ encodes the expected vertical displacement in busy hours between the post and pre COVID-19 relationships for the same traffic level. Negative values indicate that, for the same N, the post COVID-19 schedule has fewer hours above the high-utilisation threshold. To account for data variability, we normalise this as $\hat{\delta}_B = \delta_B / \sigma_p$, where σ_p is the pooled standard deviation of residuals from both fits - equivalent to Cohen's D [16]. For the synthetic example in Fig. 2, $\hat{\delta}_B \approx -1.33$, indicating a reduction in busy hours exceeding the metric's variability. The right panel of Fig. 2 reports both $\hat{\delta}_B$ and $\hat{\delta}_H$ for airports that have recovered at least 90% of pre COVID-19 traffic, ensuring sufficient overlap for reliable fits. Most airports show $\hat{\delta}_B$ values near zero, suggesting no major changes; however, three (EIDW, LIRF, EGSS) exhibit strongly negative values,

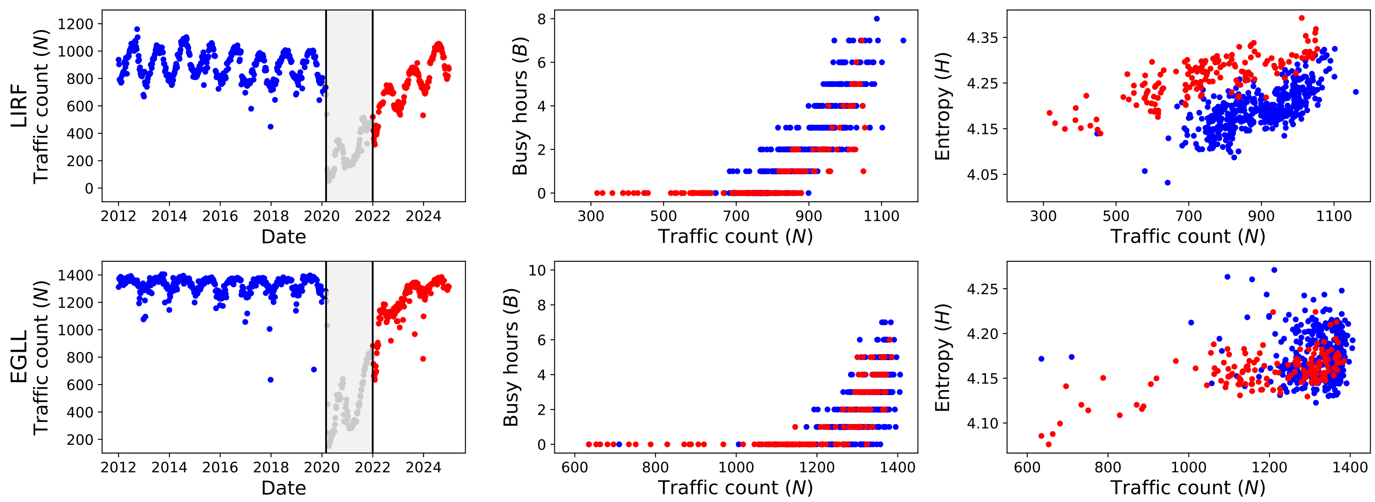

Figure 1. Characterising operations at LIRF (top panels) and EGLL (bottom panels). From left to right, panels depict: the evolution of the daily number of operations; the daily number of busy hours, as a function of the number of operations; and entropy of operations, also as a function of their number. Blue and red points respectively correspond to pre and post COVID-19 days. Changes in peakiness are assessed by comparing blue and red points vertically at similar values of N.

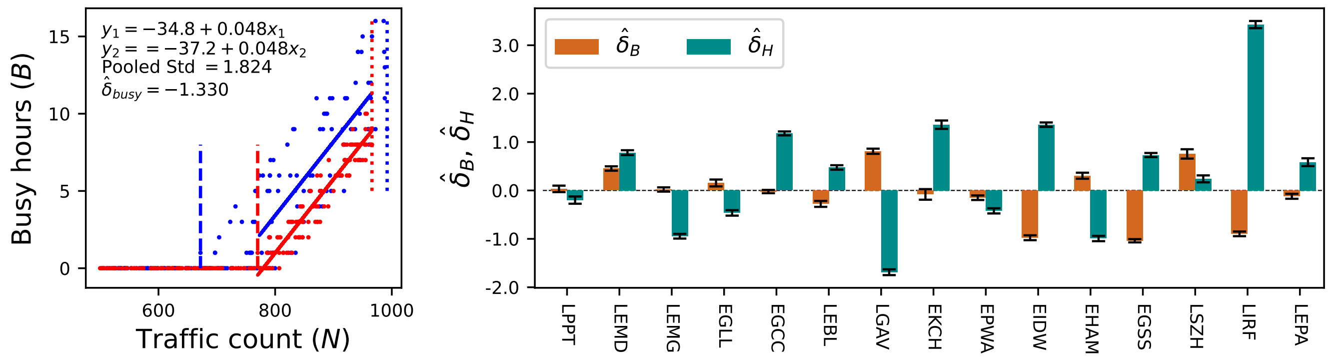

Figure 2. Relationship between dynamics with the traffic. Here, N is the daily number of operations, B is the number of busy hours, and H is the entropy of the hourly traffic profile; negative $\hat{\delta}_B$ values indicate fewer busy hours after COVID-19, whereas positive $\hat{\delta}_H$ values indicate a more depeaked post COVID-19 distribution of operations. (Left) Example of the calculation of $\hat{\delta}_B$ on synthetic data. The diagonal solid lines represent the best linear fits, whose parameters are reported in the top left part. Dashed vertical lines (left side) represent the minimum N for which $B > 0$; dotted ones (right side) the maximum value of N available for the airport. Blue and red colours correspond to pre and post COVID-19. (Right) Obtained $\hat{\delta}_B$ and $\hat{\delta}_H$ for all airports that have recovered at least 90% of the traffic. Whiskers indicate the confidence intervals, obtained using a bootstrapping that randomly samples 80% of the data without replacement, and 50 independent realisations. See main text for definitions.

while LSZH shows a strongly positive one. $\hat{\delta}_H$ is calculated analogously for entropy (see the right-hand panels of Fig. 1), using the full range of N for fitting. This represents how entropy changed between periods at comparable traffic levels, accounting for variability. Positive values indicate that post COVID-19 schedules are more uniformly distributed than pre COVID-19 ones, for the same number of daily operations. While LIRF shows a very large value, all other airports- including those with notable $\hat{\delta}_B$ changes - display more modest $\hat{\delta}_H$ values.

To facilitate interpretation, Fig. 4 reports scatter plots of $\hat{\delta}_B$ and $\hat{\delta}_H$ versus the percentage of recovered traffic. The relationship for both is weak and statistically not significant. The figure, therefore, does not indicate a general shift towards less peaked schedules. Instead, it reinforces what has been seen throughout this work: changes in schedule peakiness are heterogeneous across airports, and traffic recovery alone is insufficient to accurately explain them.

IV. Evolution through Time

Previous analyses characterised busy hours and entropy as functions of traffic volume across entire periods. We now examine their temporal evolution within each period. However, these aspects are intertwined: entropy may decrease over time simply due to increasing traffic - the system is non-stationary. To account for this, we normalise entropy by traffic volume using the period-specific fits from Sec. III to extract the expected entropy given the observed traffic:

$$H_e(N) = \beta N + \alpha_{\text{period}}. \tag{3}$$

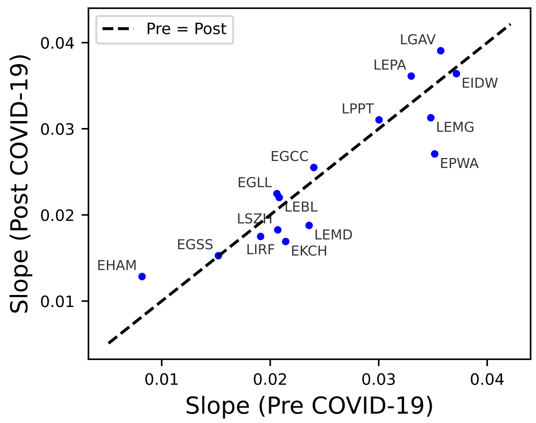

Figure 3. Comparison between the slopes obtained by fitting the relationship between busy hours B and daily traffic N independently in the pre and post COVID-19 periods. Each point corresponds to one airport. The diagonal line represents equality between the two slopes.

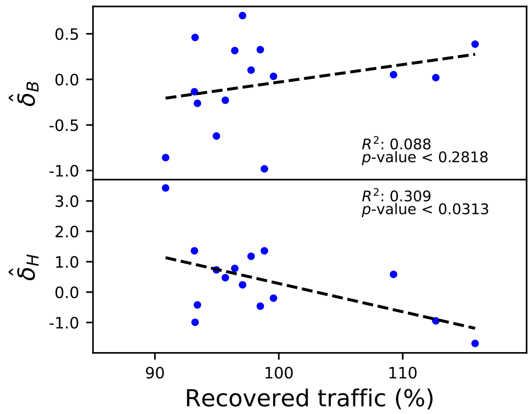

Figure 4. Analysis of the entropy and number of busy hours. From top to bottom, panels synthesise the results obtained in the right panels of Fig. 2, specifically: $\hat{\delta}_B$ as a function of the percentage of recovered traffic; and $\hat{\delta}_H$ as a function of the same metric. Dashed black lines depict the best linear fits.

For each day t, we calculate residuals - how much observed entropy deviated from expectations for that day's traffic:

$$\Delta H(t) = H(t) - H_e[N(t)], \tag{4}$$

then normalise by the standard deviation:

$$\hat{\Delta} H(t) = \frac{\Delta H(t)}{\sigma_{\Delta H}}, \tag{5}$$

yielding a dimensionless measure equivalent to Cohen's D. Trends in $\hat{\Delta} H(t)$ reveal changes in temporal organisation unexplained by traffic volume.

Fig. 5 shows the evolution of $\hat{\Delta} H$ for LIRF and EGLL. Linear fits reveal contrasting scenarios: the $\hat{\Delta} H$ value increases slightly for LIRF (especially post COVID-19), while it decreases for EGLL. Analysing the slope of the linear fits of $\hat{\Delta} H$ versus the recovered traffic for all airports reveals a weak positive correlation in the post COVID-19 period: airports exceeding 2019 traffic levels show increasing entropy, even when accounting for increased traffic.

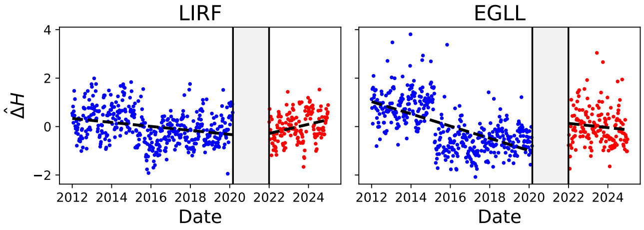

Figure 5. Evolution of the normalised entropy of both schedules, for LIRF (left) and EGLL (right); see main text for definitions. Blue and red points respectively correspond to pre and post COVID-19. Lines represent a linear fit of the data for each period.

V. Role of Leading Airlines

The airport-level analysis presented above captures the aggregate temporal load experienced at each airport, but it should not be interpreted as evidence that airports themselves determine schedule structure. Flight schedules are primarily generated by airlines, subject to operational constraints. Banked schedules may be valuable for airlines and passengers because they increase the number of feasible connections and reduce transfer times. At the same time, the aggregation of several airline schedules determines the temporal concentration faced by airport infrastructure and local operational resources. A natural extension is therefore to decompose the airport-level indicators by dominant airline, calculating entropy and $\hat{\delta}_H$ separately for each airline's operations.

Preliminary results applying this decomposition reveal heterogeneous carrier-level contributions. At LIRF, ITA Airways/Alitalia contributed substantially to depeaking ($\hat{\delta}^d_H = 2.24$), with other airlines also contributing positively ($\hat{\delta}_H = 0.96$). However, dominant airlines sometimes push against overall airport trends: Ryanair at LEMG and British Airways at EGLL showed positive entropy shifts while their airports' overall entropy decreased; conversely, Swissair at LSZH showed negative shifts while the airport's entropy increased. In some cases, secondary airlines had greater impact despite smaller market shares - for instance, Aer Lingus at EIDW ($\hat{\delta}^d_H = 1.45$) versus Ryanair ($\hat{\delta}^d_H = 0.45$). Future work will systematically quantify these airline-level contributions and explore the strategic incentives driving divergent behaviours.

VI. Discussion and Future Work

Are airport schedules in Europe becoming more or less peaked? Our analysis suggests that there is no general trend in any direction. Most airports maintained similar temporal concentration before and after the pandemic for comparable number of operations, with traffic peaks intensifying at those exceeding pre-pandemic levels. However, a few notable exceptions emerge. Rome Fiumicino (LIRF) exhibits a clear structural shift toward a more even distribution of operations, with similar but less pronounced effects observed at Dublin (EIDW) and London Stansted (EGSS). By contrast, airports such as Zurich (LSZH) show a shift in the opposite direction, while highly saturated hubs like London Heathrow (EGLL) are constrained to accommodating traffic wherever capacity allows, leaving little room for structural change. These contrasts reflect local constraints: where capacity is tight, bank patterns persist or sharpen; where slack exists, smoothing becomes feasible.

Airlines play a central role in these dynamics. Preliminary decomposition by carrier reveals heterogeneous behaviours: at LIRF multiple airlines contributed to depeaking, while elsewhere dominant carriers sometimes push against airport-wide trends. This dispersion reflects underlying network logic - hub-and-spoke carriers preserve banks for connectivity, while point-to-point operators may adopt rolling patterns depending on local constraints [17]. This airline-level heterogeneity highlights a fundamental coordination problem in schedule formation. Individual carriers optimise their own network performance - connectivity, fleet rotations, crew pairing, and passenger preferences - while the resulting temporal concentration is a collective outcome experienced at the network level. Hence, the airport-level view hides important differences in airline business models: airports dominated by hub carriers naturally show stronger peaks, while those with more low-cost or point-to-point traffic may find smoother schedules easier to achieve.

Several extensions of the current work are worth exploring. First, hourly aggregation may hide finer dynamics; analysing windows shorter than one hour could reveal patterns within peaks and better capture short-term banks. Second, while we show that peaking persists, quantifying its consequences in terms of delays, resource use, passenger connectivity and travel times would allow a more complete assessment of the trade-offs between banked and less concentrated schedules [18]–[20]. Third, a systematic decomposition of airline-level contributions, building on the preliminary results in Sec. V, would clarify whether the patterns we observe reflect airline strategy or structural limits. Finally, extending the analysis to non-European airports would help clarify whether the patterns observed here reflect global scheduling behaviours or are specific to the European context.

Acknowledgment

J.B. acknowledges economic support by the Conselleria d'Educació i Universitats of the Government of the Balearic Islands (FPI 042/2023). This work was partially supported by the project Hybrid within the Engage 2 program of the SESAR Joint Undertaking under EU Commission project (101114648). Partial funding has been received also from the María de Maeztu project CEX2021-001164-M funded by the MICIU/AEI/10.13039/501100011033, the COSASTI (PID2024-157493NB-C22) project by the MICIU/AEI/10.13039/501100011033 and Fondo Europeo de Desarrollo Regional (FEDER, UE).

References

[1] D. L. Bryan and M. E. O'kelly, "Hub-and-spoke networks in air transportation: an analytical review," *Journal of regional science*, vol. 39, no. 2, pp. 275–295, 1999.

[2] M. Bazargan, *Airline operations and scheduling*. Routledge, 2016.

[3] Y. Zhang, M. Menendez, and M. Hansen, "Analysis of de-peaking strategies implemented by american airlines: causes and effects," in *83rd Transportation Research Annual Meeting*, 2004.

[4] M. Frank, M. Mederer, B. Stolz, and T. Hanschke, "Depeaking–economic optimization of air traffic systems," *Aerospace science and technology*, vol. 9, no. 8, pp. 738–744, 2005.

[5] H. Jiang and C. Barnhart, "Dynamic airline scheduling," *Transportation Science*, vol. 43, no. 3, pp. 336–354, 2009.

[6] D. Katz and L. A. Garrow, "Depeaking schedules: Beneficial for airports and airlines?" *Transportation research record*, vol. 2325, no. 1, pp. 43–55, 2013.

[7] J. Herring, V. Lurkin, L. A. Garrow, J.-P. Clarke, and M. Bierlaire, "Airline customers' connection time preferences in domestic us markets," *Journal of Air Transport Management*, vol. 79, p. 101688, 2019.

[8] M. Luethi, B. Kisseleff, and A. Nash, "Depeaking strategies for improving airport ground operations productivity at midsize hubs," *Transportation research record*, vol. 2106, no. 1, pp. 57–65, 2009.

[9] D. S. Katz and L. A. Garrow, "Revenue and operational impacts of depeaking at us hub airports," *Journal of Air Transport Management*, vol. 34, pp. 57–64, 2014.

[10] Performance Section, DECEA and Performance Review Unit, EUROCONTROL, "Brazil/Europe Comparison of Operational ANS Performance," DECEA and EUROCONTROL, Report, Oct. 2022, 2nd ed., published Oct. 28, 2022. Chapter 5: Capacity and Throughput.

[11] M. Lohatepanont and C. Barnhart, "Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment," *Transportation science*, vol. 38, no. 1, pp. 19–32, 2004.

[12] A. E. Eltoukhy, F. T. Chan, and S. H. Chung, "Airline schedule planning: a review and future directions," *Industrial Management & Data Systems*, vol. 117, no. 6, pp. 1201–1243, 2017.

[13] Y. Xu, S. Wandelt, and X. Sun, "Airline scheduling optimization: literature review and a discussion of modelling methodologies," *Intelligent Transportation Infrastructure*, vol. 3, p. liad026, 2024.

[14] C. E. Shannon, "A mathematical theory of communication," *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.

[15] A. Lesne, "Shannon entropy: a rigorous notion at the crossroads between probability, information theory, dynamical systems and statistical physics," *Mathematical Structures in Computer Science*, vol. 24, no. 3, p. e240311, 2014.

[16] C. O. Fritz, P. E. Morris, and J. J. Richler, "Effect size estimates: current use, calculations, and interpretation." *Journal of experimental psychology: General*, vol. 141, no. 1, p. 2, 2012.

[17] S. Gil-Rodrigo and M. Zanin, "Low cost carriers induce specific and identifiable delay propagation patterns: an analysis of the eu and us systems," *IEEE Access*, vol. 12, pp. 75 323–75 336, 2024.

[18] P. Fleurquin, J. J. Ramasco, and V. M. Eguiluz, "Systemic delay propagation in the us airport network," *Scientific reports*, vol. 3, no. 1, p. 1159, 2013.

[19] A. Sismanidou, J. Tarradellas, and P. Suau-Sanchez, "The uneven geography of us air traffic delays: Quantifying the impact of connecting passengers on delay propagation," *Journal of Transport Geography*, vol. 98, p. 103260, 2022.

[20] M. Nasrollahi and A. Abdi Kordani, "Designing airline hub-and-spoke network and fleet size by a biobjective model based on passenger preferences and value of time," *Journal of Advanced Transportation*, vol. 2023, no. 1, p. 2797613, 2023.