Robustness analysis of European rail and air service networks

Robustness analysis of European rail and air service networks

Master of Science Thesis

in partial fulfilment of the requirements for the degree of

Master of Science
in Civil Engineering

by

Hanyu Cheng

MSc Civil Engineering (Track Traffic and Transport Engineering)
Department of Transport & Planning
Faculty of Civil Engineering and Geosciences
Delft University of Technology
Delft, South Holland, The Netherlands

Thesis Committee:

Chair:	**Prof. dr. Oded Cats**
Supervisor:	**Dr. Yanan Xin**
Supervisor:	**Dr. Dipl.-Ing. Rajat Verma**

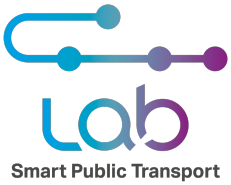

Keywords: Network robustness, Complex networks, Multilayer networks, Intermodality.

Printed by: Hanyu Cheng

Front & Back: NS trains in snow.

Contents

List of Figures

List of Tables

List of Acronyms

General acronyms

The third column points to the section in which each acronym is first used. Operator, airline, airport, and country codes that appear only as data labels in Chapter 5 and the appendix are not listed here.

Acronym	Expansion	First used
API	Application Programming Interface	Section 3.2.2
ATC	Air Traffic Control	Section 2.1.1
ATFM	Air Traffic Flow Management	Section 5.1
ATN	Air Transport Network	Section 5.1.1
CSA	Connection Scan Algorithm	Section 3.2.5
EASA	European Union Aviation Safety Agency	Section 2.1.1
EEA	European Economic Area	Section 1.3
EFTA	European Free Trade Association	Section 1.3
EMS	Emergency Management Service (Copernicus)	Section 5.3
ERA	European Union Agency for Railways	Section 2.1.1
ERTMS	European Rail Traffic Management System	Section 2.1.1
EU	European Union	Section 1.1
EUROCONTROL	European Organisation for the Safety of Air Navigation	Section 2.1.1
FUA	Functional Urban Area	Section 1.3
GDP	Gross Domestic Product	Section 2.1.1
GTC	Generalised Travel Cost	Section 2.2.3
GTFS	General Transit Feed Specification	Section 1.3
HSR	High-Speed Rail	Section 2.1.1
IATA	International Air Transport Association	Section 2.1.1
IC	Intercity	Section 3.2.2
ICE	Intercity-Express	Section 5.1.1
IVT	In-Vehicle Time	Section 3.2
LCC	Largest Connected Component (and, where context makes the distinction clear, Low-Cost Carrier)	Section 2.1.1
MITA	Multilateral Interline Traffic Agreement	Section 5.2.3
MLN	Multilayer Network	Section 1.3.2
MLPTN	Multilayer Public Transport Network	Section 1.3.2
MLTN	Multilayer Transport Network	Section 1.3.2
NUTS	Nomenclature of Territorial Units for Statistics	Section 3.1.1
OD	Origin-Destination	Section 3.1.1
OECD	Organisation for Economic Co-operation and Development	Section 3.1.1
OVT	Out-of-Vehicle Time	Section 3.2

List of Acronyms continued from previous page

Acronym	Expansion	First used
PT	Public Transport	Section 4.1
PTN	Public Transport Network	Section 1.1
RAPTOR	Round-Based Public Transit Optimised Router	Section 3.2.5
SES	Single European Sky	Section 2.1.1
TAG	Transport Analysis Guidance (UK Department for Transport)	Section 3.2.2
TEN-T	Trans-European Transport Network	Section 1.2
TGV	Train à Grande Vitesse	Section 3.1.3
TOC	Train Operating Company	Section 5.2.3
VAAC	Volcanic Ash Advisory Centre	Section 5.3.1

Country codes

Two-letter country codes follow the Eurostat convention, so the United Kingdom is UK and Greece is EL.

Code	Country	Code	Country
AT	Austria	IE	Ireland
BE	Belgium	IT	Italy
BG	Bulgaria	LT	Lithuania
CH	Switzerland	LU	Luxembourg
CY	Cyprus	LV	Latvia
CZ	Czechia	MT	Malta
DE	Germany	NL	Netherlands
DK	Denmark	NO	Norway
EE	Estonia	PL	Poland
EL	Greece	PT	Portugal
ES	Spain	RO	Romania
FI	Finland	SE	Sweden
FR	France	SI	Slovenia
HR	Croatia	SK	Slovakia
HU	Hungary	UK	United Kingdom

Rail operator codes

The supply layer is served by 51 rail operators, catalogued in full in appendix Table 6.3. Only those whose label is an abbreviation are decoded below; the rest appear under their own name. Names follow the operating-company field of the intercity route dataset.

Code	Operator
CD	České dráhy
CFL	Société Nationale des Chemins de Fer Luxembourgeois
CFR	CFR Călători (Căile Ferate Române)
CP	Comboios de Portugal
DB	Deutsche Bahn (including DB Fernverkehr)
DFTO	DfT Operator Ltd (UK government-operated services)
DSB	Danske Statsbaner
EMR	East Midlands Railway
GTR	Govia Thameslink Railway
HZZP	HŽ Putnički prijevoz (Croatian Railways)
MÁV	Magyar Államvasutak
NMBS/SNCB	Belgian Railways
NS	Nederlandse Spoorwegen
PKP	PKP Intercity
Renfe	Renfe Operadora
SBB	Schweizerische Bundesbahnen
SJ	Statens Järnvägar (SJ AB)
SNCF	SNCF Voyageurs
SOB	Schweizerische Südostbahn
VR	VR Group (Finnish Railways)
ZSSK	Železničná spoločnosť Slovensko
ÖBB	Österreichische Bundesbahnen

Air carrier codes

The 31 air carriers of the supply layer, by IATA airline designator.

Code	Carrier
A3	Aegean Airlines
AF	Air France
AY	Finnair
AZ	ITA Airways
BA	British Airways
BT	airBaltic
DY	Norwegian
EW	Eurowings
FB	Bulgaria Air
FI	Icelandair
FR	Ryanair

Air carrier codes continued from previous page

Code	Carrier
HV	Transavia
IB	Iberia
KL	KLM
KM	Air Malta
LG	Luxair
LH	Lufthansa
LO	LOT Polish Airlines
LS	Jet2
LX	Swiss
OS	Austrian Airlines
OU	Croatia Airlines
QS	Smartwings
RO	TAROM
SK	SAS
SN	Brussels Airlines
TO	Transavia France
TP	TAP Air Portugal
U2	easyJet
V7	Volotea
W6	Wizz Air

Nomenclature

The symbols below are gathered from Chapter 3 and the result chapter. Each symbol is reused with the same meaning throughout the thesis. Numerical baseline values are given where the symbol is treated as a calibrated constant. The third column points to the equation or section in which the symbol is first introduced.

Network and multilayer graph

Symbol	Description	First used
$\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L}, \mathcal{T}, w)$	Multilayer service graph used throughout the thesis.	Section 3.1
$\mathcal{V} = V^r \cup V^a$	Node set, partitioned into rail stations V^r and airports V^a.	Section 3.1
$\mathcal{E} = E^r_{\text{intra}} \cup E^a_{\text{intra}} \cup E_{\text{inter}}$	Edge set, partitioned into intra-layer service edges per mode and inter-layer footpath edges.	Section 3.1
$\mathcal{L} = \{r, a\}$	Layer index over rail r and air a.	Section 3.1
$\mathcal{T}$	Operating time index over the 24 hourly departure windows.	Section 3.1
$w : \mathcal{E} \to \mathbb{R}_{\geq 0}$	Edge-weight function returning the generalised travel cost.	Section 3.1
A^r, A^a	Per-layer service adjacency matrices for rail and air.	Section 3.1.3
C^{ra}, C^{ar}	Directional inter-layer coupling matrices for rail-to-air and air-to-rail transfers.	Section 3.1.3
$\mathbf{A} = [A^r, C^{ra}; C^{ar}, A^a]$	Supra-adjacency matrix of the multilayer graph.	equation 3.9
N, $\lvert V \rvert$	Number of nodes in the working graph.	Section 3.1
m	Total number of edges.	Section 3.1
d_{ij}	Shortest path length between nodes i and j.	Section 3.1
$N(i)$	Set of surviving neighbours of node i.	Section 3.1
N_{OD}, $\mathcal{OD}$	Number and set of origin-destination pairs in the working universe.	Section 3.1.1
dep_win $\in \{0, \ldots, 23\}$	Canonical hourly departure window assigned to every path.	Section 3.2.5

Generalised travel cost

Symbol	Description	First used
GTC	Generalised travel cost of a path, in in-vehicle-minute equivalents.	equation 3.21
$\mathrm{GTC}^{(x)}_{\min}(b)$	Minimum GTC over the surviving alternatives of bucket b at disruption fraction x.	equation 3.24
t^{ivt}_l	In-vehicle travel time of leg l, summed across all legs of the path in the first GTC term.	Section 3.2.1
$W^{\mathrm{o}}(p), W^{\mathrm{d}}(p)$	Origin and destination dwell priced at the first and last terminal of path p, read from Table 3.4 by mode and hub class.	equation 3.21
$\mathscr{T}(p), \mathscr{T}^{\mathrm{move}}(p)$	Transfer set of path p and its inter-terminal subset, with at most two transfers and at most one flight-to-flight change.	equation 3.21
t^{gap}_τ	Elapsed time between the arrival of one leg and the departure of the next at transfer τ.	Section 3.2.3
$t^{\mathrm{connect}}_\tau$	Intra-FUA connecting time between the alighting and the next boarding facility at transfer τ, zero at a same-terminal transfer.	Section 3.2.3
t^{wait}_τ	Platform-side waiting at transfer τ, the residual after subtracting the arrival buffer and the connecting time from the gap.	Section 3.2.3
$B_{\mathrm{mode}(k)}$	Arrival-side buffer at the alighting facility, a feasibility guard that is not priced, 50 min at a rail station, 90 min at a medium or small airport, 120 min at a large hub.	equation 3.17
$W^{\mathrm{mode}(k')}_{\min}$, $W^{\mathrm{mode}(k')}_{\max}$	Boarding floor and ceiling on the transfer gap at the boarding facility, 5 and 60 min at a rail station, 90 and 1,440 min at a medium or small airport, 120 and 1,440 min at a large hub.	equation 3.17
$\mu_{\mathrm{wait}} = 1.76$	Waiting-time multiplier on the endpoint dwells and all platform-side waiting, calibrated from Wardman et al. (2016).	Section 3.2.2
$\nu_{\mathrm{move}} = 1.87$	Access and egress multiplier on the intra-FUA connecting time $t^{\mathrm{connect}}_\tau$, calibrated from Wardman et al. (2016).	Section 3.2.2

Path enumeration

Symbol	Description	First used
$\mathcal{T}_r, \mathcal{T}_a$	Rail and air timetable inputs to the path-finding procedure.	Section 3.2.5
b, $\mathcal{B}$	Departure bucket, an origin-destination pair with an hourly departure window, and the set of all buckets.	Section 3.3
$F(s, s', \text{voronoi\_win})$	Intra-FUA connection lookup returning the connecting time between facilities s and s' in the specified time-of-day band.	Section 3.2.5
Π_k	Set of feasible paths with k transfers, for $k \in \{0, 1, 2\}$.	Section 3.2.5
$p \oplus \ell$	Concatenation of path p and direct leg ℓ.	Section 3.2.5
$\text{cal}(\cdot)$	Operating-date bitmask of a leg or path.	Section 3.2.5
$\Delta_{\text{band}} = 60$ min	GTC band filter, retaining only paths within Δ_{band} of the bucket minimum.	Section 3.2.5
$C_{\max} = 1{,}440$ min	Terminal-to-terminal maximum elapsed cap of a path.	Section 3.2.5

Topological metrics

Symbol	Description	First used
S	Size of the largest (giant) connected component, as a fraction of nodes.	Section 2.2.2
$E(G)$	Topological global efficiency.	equation 2.1
$\bar{k}_{\text{str}}$, $\bar{k}_{\text{sup}}$	Country-level mean degree in the structural and supply layers.	Table 4.1
$\langle k \rangle$, $k_{\max}$	Mean and maximum node degree of a layer.	Section 4.2
$\kappa = \langle k^2 \rangle / \langle k \rangle^2$	Degree heterogeneity ratio of a layer.	Section 4.2
$\hat{\alpha}$, $x_{\min}$	Fitted power-law tail exponent and lower cutoff of a degree distribution.	Table 4.3
$\langle C \rangle$, $\langle \ell \rangle$, d	Mean clustering coefficient, mean shortest-path length, and diameter, computed on the largest connected component.	Section 4.2
r	Degree assortativity coefficient of a layer.	Section 4.2
γ, α	Planar connectivity and meshedness indices of a layer.	Section 4.2

Service-level metrics and indices

Symbol	Description	First used
E_G	GTC-based global efficiency.	equation 3.24
x	fraction (share) of edges removed	equation 3.25
$R(x)$	Retained-ratio degradation curve at disruption fraction x.	equation 3.25
$\Delta\bar{D}(x)$	Mean additional travel time across surviving OD pairs.	equation 3.26
$f_{\text{disc}}(x)$	Fraction of OD pairs disconnected at disruption fraction x.	Section 3.3.1
$f_{\text{eff-disc}}(x)$	Fraction of OD pairs disconnected or delayed beyond the reroute tolerance at disruption fraction x.	Section 3.3.1
F	Cumulative robustness loss, the area above $R(x)$.	equation 3.28
Q	Accumulated delay, the area under $\Delta\bar{D}(x)$.	equation 3.29

Disruption scenarios

Symbol	Description	First used
Φ_s	Disruption rule returning the degraded network for an intensity and targeting-rule pair.	equation 3.30
$\mathcal{S}$	Catalogue of targeting rules.	equation 3.31
$\mathcal{U}$, $\mathcal{U}^{\text{edge}}$	Targeted element universe, and its long-distance edge form.	equation 3.32
c_s	Centrality score ranking the element universe under rule s, with removals taken as the top $K(x) = \lceil x \lvert \mathcal{U} \rvert \rceil$ elements.	equation 3.34
$b^{\text{topo}}(e)$	Topological edge betweenness on the consecutive-stop graph, the mode-blind removal ordering.	equation 3.35
$\hat{b}(e)$	Service-aware betweenness, the number of origin-destination-departure buckets whose minimum-GTC itinerary traverses edge e.	equation 3.36
$R_s(x)$	Set removed under rule s at intensity x.	equation 3.33
$G_s(x)$	Degraded network after removal, $G_s(x) = G \setminus R_s(x)$.	equation 3.33

Summary

European long-distance passenger travel is carried largely by two networks, a nationally fragmented rail system and a partially liberalised air system. The robustness of both is conventionally assessed on a topological representation, as the capacity of the graph to maintain connectivity under the removal of nodes and links. Such an assessment omits the service-level characteristics on which passengers experience a disruption, the timetable that determines when a connection exists, the transfer window that determines whether it can be used, and the operator that decides whether a stranded passenger may be rebooked. A robustness analysis of the European rail and air service networks is therefore conducted to determine to what extent these networks are robust to disruptions when service-level characteristics are incorporated into their network representation.

To answer this question two complementary multilayer transport network representations are constructed. The structural representation keeps rail and air as coupled but distinct layers whose edges are the segments between consecutive stations, the non-stop flights, and the public-transport links among the terminals of the same urban area, while the functional representation uses the structural network as its topological layer and extends it with the scheduled services that realise each connection, carrying their operator, timetable, and generalised travel cost. The two are deliberately kept apart rather than merged into one graph, so that a disruption can be aimed at the structural or the supply level independently while the shared node set keeps the two sets of results comparable. On the functional network a path-enumeration algorithm constructs the 482,947,150 feasible itineraries between 1,674 rail stations and 197 airports in 30 countries, and robustness is measured mainly through the retained efficiency $R(x)$ and the share of baseline itinerary quality that survives a disruption. Three experiment families are evaluated. Structural stress tests remove edges by unweighted topological betweenness, by a service-aware betweenness set by the itineraries each edge carries, or at random. Operator withdrawals and cooperation-tier restrictions act on the 82 rail and air operators. Two documented macro-shocks, the Iceland 2010 volcanic eruption and the July 2021 Bernd flood, test the network's response to real hazards.

Incorporating service-level information does not merely lower the measured robustness, it changes which edges are deemed critical. The two rankings share not a single rail edge among their top fifty, since topology elevates the international cut-edges that keep the continental graph connected while service load concentrates 88 per cent of the top two hundred edges on the German Rhine-Ruhr and Frankfurt-Mannheim corridors. Removing the top five per cent of edges leaves retained efficiency at $R(x) = 0.168$ in service-aware order against 0.557 in topology order, and against the random benchmark the topology failure is statistically indistinguishable ($z_R = -0.47$) while the service-aware failure is an extreme outlier ($z_R = -8.51$). The damage falls as outright disconnection rather than delay, so the service layer collapses long before the largest connected component does. Rail-air co-existence does not offset this concentration. The robustness of the hub-and-spoke air

layer does not transfer to the combined network, since an air leg reaches it through rail-side airport-feeder edges whose removal severs the intermodal itinerary at its rail end. The two macro-shocks test this in both directions. An ash closure grounding 77.7 per cent of active airports still leaves $R(x) = 0.875$ through the immunity of the rail-only baselines, with not a single air baseline replaced by a within-tolerance rail substitute. The Bernd flood lowers intra-German retained efficiency to $R(x) = 0.510$, with approximately 81,000 structurally feasible air substitutes arriving outside the reroute tolerance. In neither direction does cross-modal rescue materialise.

In sum, robustness is not a fixed property of the European air and rail networks. It depends on how the network is represented, on how severe and how targeted the disruption is, and on how service is organised across operators, from the concentration of supply in a few of them to the cooperation rules that decide whose services a passenger may combine. With service-level characteristics incorporated, they prove substantially less robust to disruptions than their topology suggests.

Preface

I've always been drawn to how varied Europe's rail and air networks are. They remind me a little of China's Warring States period, when one continent held a mix of gauges and road systems. But once you lay service information over the infrastructure, you start to see how much really comes down to timetables, transfers, and tickets. The gap between the networks we model and the journeys people actually could take is what this work has tried to close. I also didn't expect the gap to be as wide as it was, or for the binding constraint on European robustness to turn out to be more about operations management, sometimes even outweighing the pure infrastructure. Perhaps this also explains why travelling between major cities, say Copenhagen and Stockholm, or Luxembourg and Rotterdam, takes so much longer by rail than covering the same distance would in modern China.

I owe my deepest thanks to my chair, Prof. Oded Cats. Your courses, Transport Modelling and Analysis (CIEM6000) and Public Transport Demand and Network Planning and Operations (Unit 2 of CIEM6230), first drew me toward this field and shaped how I have come to think about it. Your invitation to the symposium, the core meeting with the 3mars group, your guidance on the overall track, and your detailed feedback were all invaluable. I am also deeply grateful to Dr. Yanan Xin for her patience and close attention to the methodology, and to Dr. Rajat Verma for his care and earlier help on complex networks at the East Asian urban scale. Working within the Department of Transport & Planning and the TTE track has been a privilege, and I thank the wider group for making it a place I wanted to spend my time. Looking back, I came to Delft from structural engineering, but eventually I spent two years working on transport, which is my true passion.

To my family, thank you for supporting me across the distance these two years put between us. This work is as much yours as mine. To my friends and fellow students, thank you for the company.

Use of Generative AI

Following the TU Delft OPEN Publishing policies and the COPE guidelines, the author discloses that generative AI supported parts of this work. Claude (Anthropic) was used to polish wording, and to review and check code. Nothing it produced entered the thesis unexamined since the author reviewed, edited, and verified each contribution. The ideas, the analysis, and the conclusions, and the responsibility that comes with them, remain the author's alone.

Hanyu Cheng
Delft, July 2026

1

Introduction

1.1. Background

Long-distance passenger travel has grown substantially over the past century (World Bank, 2020) and carries significant economic and societal weight by enabling the maintenance of social ties across large distances and supporting tourism and business travel (Aultman-Hall & Ullman, 2020; Malichová et al., 2022). The same pattern is also observed in Europe, where both air and rail passenger volumes have expanded steadily over the past two decades (Figure 1.1). Across European passenger transport more broadly, the two modes already play prominent roles with rail's modal share continuing to grow between 2008 and 2023 (Figure 1.2) and air transport often accounting for the largest share of passenger-kilometres after private cars across the 27 European Union (EU) member states (Figure 1.3). For long-distance travel specifically, this prominence becomes even more pronounced as rail and air together also carry a substantial share of the long-distance passenger market (Brons et al., 2023; Eurostat, 2025a; Grolle et al., 2024) and form the dominant non-car alternatives for intercity and cross-border trips above the 100 km crow-fly threshold that European transport research conventionally uses to delimit long-distance travel (Aultman-Hall & Ullman, 2020; Frei et al., 2010; Van Goeverden, 2009).

However, the transport networks on which long-distance journeys take place, and on which air and rail services operate, are sensitive to various incidents, technical failures, disruptions, extreme weather, natural disasters and antagonistic events. The consequences of such disruptions include increased travel times, cancelled trips, and social and direct economic costs (Mattsson & Jenelius, 2015). Consequently, different types of failures require tailored methods for anticipation, prevention, mitigation and restoration. Evaluating disaster impacts on transport systems is a critical first step in building a resilient network. Robustness, defined as the ability of a system to remain operational and maintain a meaningful level of functionality under disturbance, serves as the principal performance metric of this evaluation (Faturechi & Miller-Hooks, 2015). In the absence of empirical evidence and comprehensive models for network performance evaluation, infrastructure appraisals and policy instruments risk being grounded in weak assumptions, leading to costly or even counter-productive interventions (Cats, 2025).

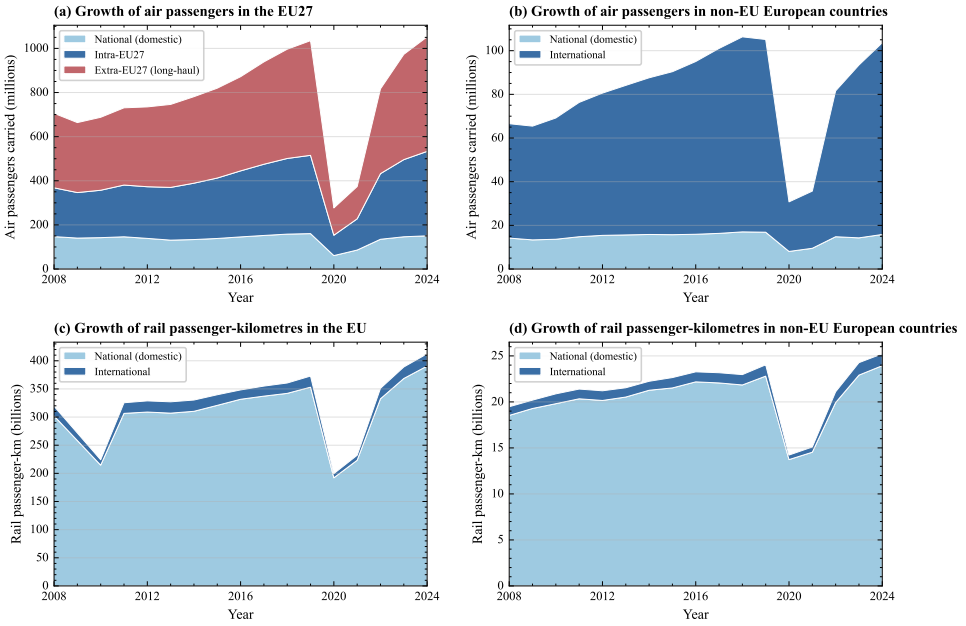

Figure 1.1: Growth of long-distance passenger travel in Europe, 2008–2024. Air panels show passengers carried and rail panels show passenger-kilometres. The EU air aggregate uses the Eurostat reporter (de-duplicated for intra-EU flows) and is split into national, intra-EU27 and extra-EU27 (long-haul) components; the non-EU European group covers Switzerland, Iceland and Norway for air and Switzerland and Norway for rail. The EU rail aggregate sums the 22 member states with continuous reporting. Sources: Eurostat (2025b, 2025c).

A coherent methodological basis is therefore needed for defining, quantifying, and benchmarking robustness metrics across large-scale rail and multimodal networks. Such an approach has long been available in the form of complex-network theory, which has previously been successfully applied to public transport networks (PTNs) whose structure involves interactions between infrastructure layout, route design, and service organisation (Luo et al., 2019; von Ferber et al., 2009). By representing heterogeneous transport systems as standardised, abstracted graphs, complex-network theory enables the derivation of topological indicators that are comparable across networks differing in infrastructure design, service characteristics, and spatial coverage. Within this framework, robustness can be more precisely defined in topological terms grounded in network science theory as the capacity of a network to maintain global connectivity and communication performance under the removal of nodes or links, whether through random failure or deliberate attack (Albert et al., 2000). More recently, the same theoretical framework has been extended to large-scale continental transport systems, despite their differences from PTNs in infrastructure layout, route design, service organisation, and service frequency. It has been applied to the European international railway network (Calzada-Infante et al., 2020) and to the European air traffic network (Cardillo et al., 2013; Hackett et al., 2016; Pien et al., 2015; Zhou et al., 2019), as well as jointly to a multimodal characterisation of 124 European metropolitan areas (Ippolito & Cats, 2024).

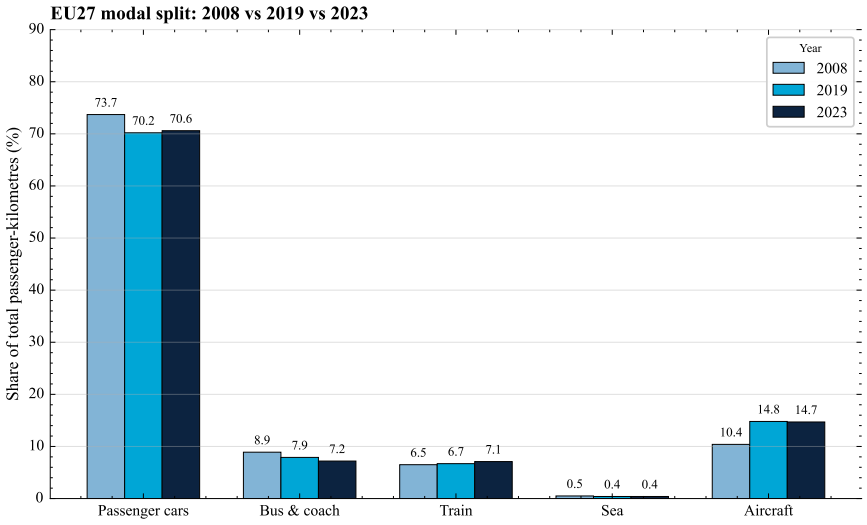

Figure 1.2: Modal split of passenger transport in the European Union, reported for the current 27 Member States, in 2008, 2019 and 2023. Values are expressed as shares of total passenger-kilometres across air, sea and inland passenger transport. Source:Eurostat (2025a).

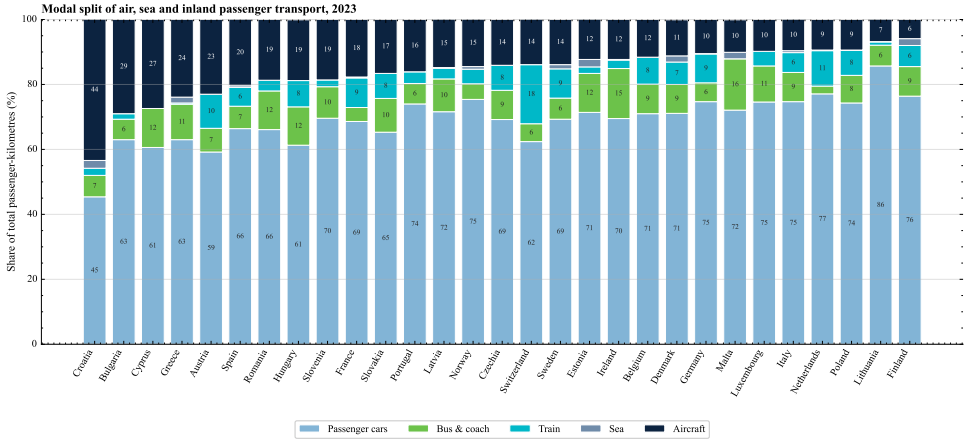

Figure 1.3: Modal split of air, sea and inland passenger transport by country, 2023. Countries are sorted by aircraft share (descending), illustrating the heterogeneity of long-distance reliance across Europe. Source: Eurostat (2025a).

Despite these advances, long-distance travel remains comparatively under-studied because of its non-routine nature, low frequency, and the difficulty of collecting representative data (Aultman-Hall & Ullman, 2020; Ullman & Aultman-Hall, 2020). This broader imbalance is also reflected in robustness research grounded in network science, which has focused predominantly on urban and regional transport systems. As a result, its application to long-distance, multimodal passenger networks remains limited, despite such

systems are more appropriately understood as network-of-networks (Cats, 2025). A similar pattern of research prioritisation can also be observed in Europe, where long-distance passenger travel has historically received limited attention in both national and supranational research programmes, partly because research commissioning bodies tend to prioritise problems within their own jurisdictions (Van Goeverden, 2009). As a result, the robustness of the European long-distance network at the continental scale has so far received insufficient systematic attention (Cats, 2025). This neglect is consequential because long-distance trips are sensitive to macro shocks such as pandemics, geopolitical events, and fuel-price fluctuations, meaning that their vulnerability is structurally different from that of daily urban trips (Malichová et al., 2022). If their robustness continues to be overlooked, infrastructure planning and resilience strategies will remain misaligned with a significant share of continental mobility flows.

Furthermore, among recent studies extending complex-network theory to Europe-wide transport systems, most still rely predominantly on purely topological representations. For example, robustness analyses such as Ippolito and Cats (2024) and Pien et al. (2015) construct directed graphs in which edges represent the existence of services between nodes derived from observed routes or flight profiles. In these representations, edges are treated as binary relationships indicating whether a connection exists. Consequently, the resulting networks remain static and retain only limited temporal or operational information about services and therefore they are insufficient to capture service dynamics, passenger impacts, or recovery actions. Relatedly, as reviewed by Bešinović (2020) on rail networks, topological analyses grounded in complex-network and graph-theoretic traditions assess network structure under the assumed failure of specific components, while often disregarding the dynamic effects of disruption on service performance. A similar limitation arises in air-network studies as Cardillo et al. (2013) demonstrate, the single-graph projection of the European multi-airline air network systematically overstates resilience by collapsing the layer-wise structure of carrier sub-networks, thereby concealing the within-layer rerouting and seat-capacity constraints that passengers actually face when an airline's flight is cancelled.

A further concern is that purely topological representations may distort the effective connectivity of the network, since when operational aspects are considered, the importance of some edges and nodes shifts, so that stations critical for operations are not necessarily also critical for the infrastructural network (Bešinović, 2020). Consequently, Ippolito and Cats (2024) and Bešinović (2020) explicitly stress the need for future research to extend such static graph representations by incorporating time-dependent graphs accounting for temporal service dynamics. This thesis directly addresses that gap, by constructing a service-aware, timetable-grounded representation of the European rail and air networks (Chapter 3) and evaluating their responses to different failure under both unimodal and intermodal scenarios (Chapter 5).

This general limitation is especially pronounced in the European context. For instance, on the rail network, the gap between infrastructural and operational connectivity is reinforced not only by physical bottlenecks, but also by institutional and regulatory fragmentation associated with multi-operator governance (Beria et al., 2023; Cats, 2025), interoperability gaps across national signalling, gauge, and power systems (Bairras & Aguas Ardaiz, 2025), and the lack of integrated ticketing and through-booking across national

systems (Raad voor de Leefomgeving en Infrastructuur, 2020). Hence, it can be concluded that the interconnectivity of European rail networks, and of long-distance transport networks more generally, is not determined by physical infrastructure alone. Rather, network connectivity is shaped and constrained by the extent to which capacity is used efficiently, technical interoperability is achieved, timetables are coherently coordinated, and digital systems are integrated (Platform for International Rail Passenger Transport, 2024). Topological analysis alone is therefore insufficient to capture the true level of connectivity experienced by passengers (Toet et al., 2026).

Viewed in combination, the observations above reveal a clear research gap. The growing importance of European long-distance passenger travel has not been matched by a robustness literature capable of capturing these networks at the level of realism their operation requires. Existing studies remain predominantly urban in focus, topological in representation, and rarely integrate the topological and operational dimensions needed to analyse the continental-scale robustness of rail and air networks jointly. A more realistic Europe-wide robustness perspective is therefore needed, one that focuses on long-distance rail and air transport and represents connectivity in a way that more closely reflects actual service provision. In response to this gap, this thesis develops a perspective by constructing a multilayer, service-aware representation of the European rail and air networks, in which edges carry information on routes, timetables, and generalised travel costs. It evaluates robustness under random failures, targeted attacks, and disruption scenarios using passenger-experience-related metrics rather than topology alone, and analyses the two modes both individually and jointly in order to assess whether intermodal integration enhances network robustness. The remainder of this chapter sets out the motivation, scope, knowledge gaps, and research questions that structure the study.

1.2. Motivation

The development of the more comprehensive perspective that incorporates operational realities, passenger experience, and scalable complex-network representations at the Europe-wide level is motivated by several concrete considerations that extend beyond addressing a methodological gap. Such a perspective has implications for the direction of European transport policy, for the specific character of long-distance travel, and for the interdependence of rail and air at both the service and network levels.

First and foremost, Europe is pursuing transport integration at higher level through initiatives such as the Trans-European Transport Network (TEN-T) (European Commission, 2022), the Single European Railway Area, and the Single European Sky (Pellegrini & Rodriguez, 2013). Despite these initiatives, both the rail and air networks remain fragmented in their current state, albeit in different ways as rail through persistent interoperability constraints and institutional barriers between existing national incumbent operators (Bairras & Aguas Ardaiz, 2025), and air through a market-structural divide between legacy hub-and-spoke carriers and low-cost point-to-point networks (Dobruszkes, 2011). Examining the robustness of both networks in their present state is nevertheless valuable, as it helps identify the vulnerabilities that may be inherited by future integrated infrastructure and service arrangements, thereby guiding the design of integration strategies. By analysing the European transport system as a continental network rather than as a set of national systems in isolation, the thesis also aligns with the prevailing direction of

European transport policy, in which cross-border integration, rather than national optimisation, has become the dominant planning paradigm (Cats, 2025).

A second motivation follows directly from the subject of this thesis, the Europe-wide rail and air networks that make long-distance travel possible. As set out in Section 1.1, these networks carry much of Europe's intercity and cross-border long-distance travel, which has grown steadily over the past two decades, yet this travel remains relatively under-researched compared with the extensive literature on urban public transport networks. Although long-distance trips occur less frequently than daily journeys, they account for a disproportionately large share of transport energy use and emissions (Malichová et al., 2022). They also play an important social and economic role by sustaining social ties across large distances, supporting tourism, and facilitating business and economic exchange. Prolonged disruption at the national and international scale can therefore erode social capital and reduce quality of life (Aultman-Hall & Ullman, 2020). It is consequently necessary for operators and infrastructure managers to understand network vulnerability and robustness so that mitigation strategies can be designed to minimise such losses.

The third motivation concerns the operational interdependence of rail and air under modal-shift policy. Constraints or interventions placed on one mode, such as short-haul flight bans or competition-driven service withdrawals, typically generate large fluctuation in demand on the alternative mode, and it is therefore necessary to assess whether the complementary mode has sufficient service-layer capacity to absorb such fluctuations. As Bruno et al. (2025) argue, operational adjustments in support of air-to-rail modal-shift policies cannot be evaluated in isolation from rail-system capacity and network design. This interdependence concerns studying rail and air in combination rather than in isolation, since the robustness of the European long-distance system depends on how the two modes interact.

A further motivation is that the study strengthens the conceptual framework for intermodal air and rail networks with a data-driven methodology. Cats (2025) identifies disaggregate route-itinerary analysis for long-distance rail and integrated air and rail services as almost non-existent in the existing literature, and lists it among the open research challenges of the long-distance transport market. By comparing route itineraries in unimodal networks with those in intermodal networks, where airlines may use rail services as spokes within a hub-and-spoke system (Givoni & Banister, 2006), the study characterises the air and rail relationship at both the aggregate and the corridor level. This makes it possible to identify the corridors and city pairs on which competition, cooperation, or integration is most likely to dominate, rather than treating substitution as a uniform outcome across the whole network (F. Zhang et al., 2018).

Finally, although the study does not explicitly model passenger flows, the network representation it constructs can serve as a first step towards the kind of long-distance passenger assignment models that Cats (2025) identifies as an unresolved challenge. This challenge, alongside the route-choice gap discussed above, arises from the fact that existing assignment models have generally been developed for either high-speed rail or airline networks in isolation. The network representation also lays a foundation for integrated air–rail assignment modelling by constructing a service-weighted graph on which such an assignment can subsequently be built, since passenger on-network times are already represented in the generalised travel-cost formulation. The resulting robustness analysis also

carries direct implications for network design and operations by identifying critical nodes, vulnerable corridors, and temporal bottlenecks whose protection should be prioritised in infrastructure investment and contingency planning (Jiao et al., 2020).

Together, these motivations demonstrate the necessity of the present study, which not only addresses a methodological gap in the literature but also has direct relevance for the policy, planning, and design of the European long-distance passenger transport system in the near future.

1.3. Scope and delimitations

1.3.1. Geographical scope

The geographical scope of this study is delimited along three dimensions, namely geographical definition of Europe, participation in the TEN-T initiatives, and the extent of transport service supply represented. The TEN-T framework is adopted as the organising boundary because it expresses the direction in which European transport is integrating, as argued in Section 1.2. Scoping the study to the countries it covers therefore aligns the analysis with the continental network that the EU is actively working towards rather than with the set of national systems as they currently stand in isolation.

The geographical definition concerns the boundaries of Europe and determines what is regarded as a **European** air or railway link. Unlike the Atlantic, Northern European, and Mediterranean boundaries, Europe's eastern borders with Turkey, Russia, and the trans-Caucasian republics lack a clear definition (European Environment Agency, 1995). The first delimitation step therefore concerns countries located along these eastern borders. Taking into account the key corridors of the TEN-T network and the current political context, Russia and Belarus are excluded in line with the 2022 amendment to the TEN-T Regulation, under which both countries were removed from the TEN-T maps (European Commission, 2022). Despite being included in the TEN-T Eastern Partnership Transport Network (European Commission, 2024), Turkey and the three South Caucasus countries are also excluded, since their rail links with the rest of Europe remain sparse, and they lack GTFS and airline schedule data of comparable quality to those available in EU and EEA countries.

With the eastern-border exclusions in place, all remaining countries within the eastern, Atlantic, Northern European, and Mediterranean boundaries are, in principle, included. This geographically defined set is then refined into the targeted set of this study, which serves as the broader reference set before any later reduction based on the availability of air and rail data. This refinement proceeds in two stages. First, the set is filtered with reference to countries participating in TEN-T policy. Second, the resulting list is cross-checked against the country set used in the closest prior study.

The targeted set of countries in the study should cover every country to which the trans-European transport network (TEN-T) framework or its formal extensions currently apply as seen in regulation set by European Parliament and Council of the European Union (2024b). This includes (i) the 27 EU member states, (ii) the EEA members Norway and Iceland, (iii) Switzerland, an EFTA state whose land-transport agreement with the EU aligns its rail and road systems with TEN-T, (iv) the United Kingdom, retained despite its post-Brexit departure from the formal framework. The Union Connectivity Review

notes that key parts of the UK network were previously part of TEN-T and that its connectivity objectives align closely with those of TEN-T (Department for Transport (United Kingdom), 2021), (v) the six Western Balkans candidate states (Albania, Bosnia and Herzegovina, Kosovo, North Macedonia, Montenegro, Serbia), formally embedded in the TEN-T framework through the Treaty establishing the Transport Community signed in Trieste on 12 July 2017 (European Union and Western Balkans Six, 2017), with the 2024 TEN-T revision (European Parliament and Council of the European Union, 2024a) formalising this extension as the Western Balkans, Eastern Mediterranean Core Network Corridor, and (vi) Ukraine, covered by the Eastern Partnership Transport Network of the TEN-T. Additionally, five European microstates (Andorra, Liechtenstein, Monaco, San Marino, and Vatican City) lie within the geographical scope but are excluded from the targeted set due to the absence of commercially relevant long-distance air and rail services or their non-participation in the Trans-European Transport Network (TEN-T).

As a sanity check, the targeted set is also cross-checked against the country basis of the closest prior robustness study of European intermodal rail-and-air networks, Ippolito and Cats (2024), whose underlying open dataset (Tanner & Provoost, 2023) covers 112 cities in 34 countries. The present targeted set is a proper superset of the country basis used in that study, since all 34 countries covered there are likewise included in this thesis. This cross-check therefore confirms that the targeted set adopted here is both aligned with the legislative framework of TEN-T and empirically continuous with the closest prior work. The targeted country set is shown in Table 1.1.

On the basis of the list of targeted countries, the final set of countries included in the analysis is derived from this targeted set of 39 by excluding those for which neither air nor rail data are available. In line with Ippolito and Cats (2024), rail stations and airports are spatially anchored to cities of interest, represented here by Functional Urban Areas (FUAs) as defined by the Eurostat Urban Audit framework, which delineates cities together with their commuting zones across European countries. Iceland is entirely absent from the Eurostat FUA dataset and therefore cannot host any network node. For the remaining countries, inclusion requires the presence of at least one FUA containing one or more rail stations or airports supported by both route-level and trip-level data. A rail station is considered active only if it appears on at least one rail-mode route in the intercity route dataset and is supported by trip-level timetable records in the corresponding GTFS compilation. Similarly, an airport is considered active only if it appears as the origin or destination of at least one scheduled flight in the March 2025 airline schedule snapshot. Albania, Bosnia and Herzegovina, Kosovo, Montenegro, North Macedonia, and Serbia have neither GTFS rail feeds nor scheduled flights in the input data, while Ukraine and Moldova contain route or airport listings but lack the corresponding trip-level records in the dataset of scheduled airline schedule. Together with Iceland, these eight countries constitute the nine for which no valid network can be constructed within the country.

After the exclusion of these nine countries, the final set comprises 30 countries (Figure 1.4), 387 Functional Urban Areas, 1,674 active intercity rail stations, and 197 active scheduled airports. The filtering method used to determine active nodes is described in detail in Chapter 3, and the resulting per-country coverage is presented in Chapter 5. It should be noted that the analysis is restricted to intra-European trips, defined as trips whose origin and destination both lie within the 30 countries listed in Table 1.1. Trips with

Table 1.1: Targeted country set of this study (39 countries), grouped by political and geographical category and anchored on the TEN-T legal perimeter set out in Regulation (EU) 2024/1679 (European Parliament and Council of the European Union, 2024b).

Group	Countries	*n*
EU member states	Austria, Belgium, Bulgaria†, Croatia, Cyprus†, Czechia, Denmark, Estonia, Finland, France, Germany, Greece†, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta†, Netherlands, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden	27
EEA (non-EU) & TEN-T extension	Iceland*, Norway	2
EFTA (non-EEA) & TEN-T extension	Switzerland	1
Outside TEN-T, retained	United Kingdom	1
Western Balkans candidates	Albania*, Bosnia and Herzegovina*, Kosovo*, Montenegro*, North Macedonia*, Serbia*	6
Eastern Partnership	Moldova*, Ukraine*	2
Included / Total		**30 / 39**

† marks countries covered on the air layer only, either because no passenger railway exists (Cyprus and Malta) or because no GTFS rail data are available (Bulgaria and Greece), while * marks countries included in the targeted set on legal TEN-T grounds but not eventually included in this study due to the absence of a Eurostat Functional Urban Area (Iceland) or of usable air and rail data.

at least one endpoint outside this set are excluded from consideration. Within Europe, the analysis is restricted to inter-city long-distance journeys, that is, journeys whose origin and destination lie in different distinct FUAs. This filter excludes intra-urban commuting and short-distance regional trips, consistent with the long-distance focus of the thesis introduced in Section 1.1, and aligns the operational definition of long-distance with the Eurostat FUA framework rather than imposing a single hard distance threshold.

1.3.2. Network model and disruption scope

Having defined the geographical boundary of the study area, this section specifies how the real-world European rail and air system is abstracted into a network model, and how disruptions are represented on that model. The study adopts a multilayer-network representation in the formalisation of Kivelä et al. (2014) and Boccaletti et al. (2014). A multilayer network generalises an ordinary graph by partitioning its nodes and edges into a finite set of interconnected layers. Each layer is itself an ordinary graph carrying its own intra-layer edges, and nodes belonging to different layers can be joined by inter-layer edges, also referred to as couplings. Different layers may contain different node subsets, different edge definitions, and different edge weights, while still forming a unified graph-theoretic representation. This generalisation is appropriate for systems in which several types of entities and several types of relationships coexist.

Building on this general construction, this study adopts the taxonomy proposed by Šfil-

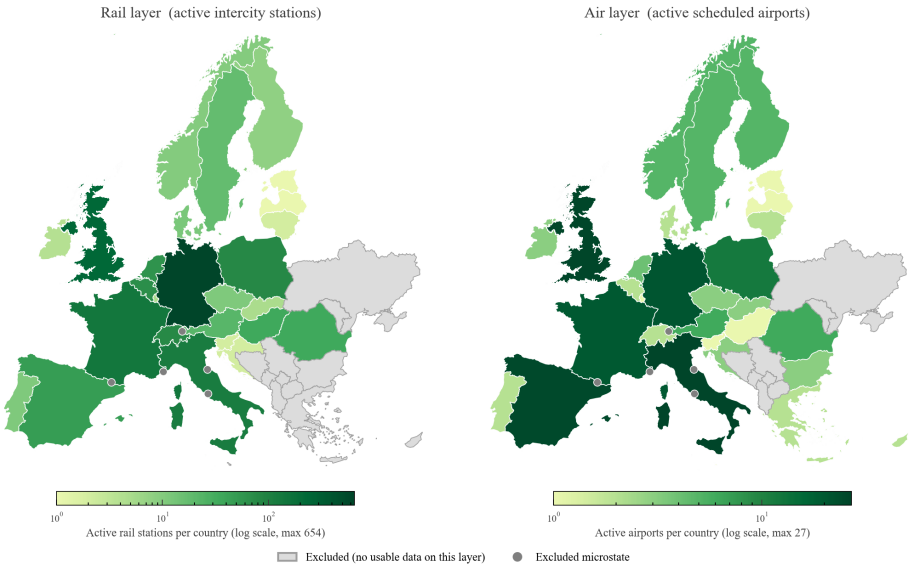

Figure 1.4: The left panel shows the number of active intercity rail stations (i.e. stations appearing in the GTFS timetable) per country; the right panel shows the number of active scheduled airports (i.e. airports with at least one scheduled commercial service in the dataset) per country. Both panels use a logarithmic colour scale. Countries shaded in grey are excluded from the respective layer for lack of usable data, and grey markers indicate excluded microstates. Iceland, excluded for the absence of a Functional Urban Area, lies beyond the north-western edge of the map and is not drawn.

igoj et al. (2026), who survey applications of multilayer-network theory to urban public transport, and extends that taxonomy from the urban public-transport context to Europe-wide rail and air transport. The resulting analytical object is therefore referred to here as a **multilayer transport network (MLTN)**. The taxonomy of Šfiligoj et al. (2026) characterises multilayer transport representations along two complementary axes. The first axis concerns the graph-theoretic relationship between layers. Layers can contain replicas of the same physical entities coupled by node identity (multiplex), conceptually similar but physically distinct entities coupled by quantitative dependency relations (interdependent), or qualitatively different entity types (heterogeneous). The second axis concerns the transport-specific basis on which layers are separated. A structural representation separates layers by the supply structure of the transport system, for example by mode or by line. A functional representation separates layers by an operational, demand, or other functional attribute overlaid on the same physical system.

Within this framework, two complementary MLTN representations are constructed and analysed in parallel. The **structural MLTN** is a mode-layered interdependent representation in which the rail layer and the air layer represent two distinct transport modes with disjoint node sets, and are coupled by inter-layer transfer edges where a rail station and an airport share a Functional Urban Area and an empirically measured walking or short urban-transit footpath connects them. The **functional MLTN** is a supply-layered representation built on the same underlying physical entities, in which the struc-

tural layer described above is paired with an operational layer that encodes scheduled timetable, operating-day calendar, and generalised travel cost on each service edge. In the two-axis taxonomy of Šfiligoj et al. (2026), the Structural MLTN sits in the mode-layered structural sub-class of the interdependent MLN type, and the Functional MLTN sits in the supply-layered sub-class of the functional MLPTN class. Running matched disruption events on both representations allows the study to compare topological robustness, read off the Structural MLTN, against service-aware robustness, read off the Functional MLTN, in line with the topology-versus-operations divergence flagged in the rail context by Bešinović (2020), in the European air context by Cardillo et al. (2013), and in the European intermodal context by Ippolito and Cats (2024).

1.4. Research questions

As mentioned in Section 1.1, two main knowledge gaps motivate this study. First, existing robustness analyses focus almost exclusively on the topological layer. While studying infrastructure unavailability is valuable in itself, this approach does not account for the service layer superimposed on the infrastructure and therefore provides no insight into the deterioration of operational and service quality following a disruption. Second, continental-scale robustness studies of European passenger networks rarely incorporate generalised travel cost or transfer penalties since edges are treated as binary, and the experienced cost of rerouting passengers after a disruption is therefore not captured. In response to these gaps, the research questions are formulated in the following Section 1.4.1 and Section 1.4.2.

1.4.1. Main research questions

The main research question of this thesis is:

> **To what extent are European air and rail networks robust to disruptions when service-level characteristics are incorporated into their network representation?**

1.4.2. Sub-questions

To address the main research question, four sub-questions are formulated. These questions correspond to the core components of the study: (i) the multilayer representation of the European air and rail networks, (ii) separate robustness analyses of the air and rail modes, (iii) robustness when the rail and air networks are combined, and (iv) implications for planning and policy.

Network representation

- How can the European long-distance air and rail networks be represented as a multilayer graph whose links carry both infrastructure and service-layer information, namely the operating route or flight, the ordered sequence of stations or airports it serves, the scheduled departure and arrival times, and the in-vehicle running time?

Unimodal service robustness

- How robust are the European air and rail service networks under the random and targeted failure of their service links, namely the rail segments between consecutive stations and the non-stop flights that make up the supply layer? When the resulting service degradation is evaluated with service-level information, how do the results compare with topology-based robustness measures?

Rail–air co-existence

- To what extent does rail–air co-existence enhance robustness at both the infrastructure and service layer, and which service links and corridors benefit most in terms of maintained connectivity under disruption?

Planning and policy implications

- Which service links and corridors are most critical for the robustness of the air network, the rail network, and the air–rail intermodal network respectively, and what targeted service or policy measures could mitigate the vulnerabilities identified?

1.5. Main contribution

The main contributions of this study are threefold.

First, the study integrates heterogeneous open data on the two modes into a single continental analytical basis. National GTFS rail feeds and commercial airline schedules differ in station and airport naming, stop-location precision, and routing conventions, and the processing pipeline harmonises them into one service-aware multilayer network of 1,674 active intercity rail stations and 197 scheduled airports across 30 countries, mapped to Functional Urban Areas. The same basis extends readily to related complex-network quantities such as accessibility and the temporal evolution of the network.

Second, as stated in the Section 1.3.2, the study designs a dual representation that holds the same European system at once as a topological object and as a service-itinerary object. The structural multilayer network carries the infrastructure layer, while the functional multilayer network overlays the timetable, the operating calendar, and a generalised travel cost on every scheduled service, with the cost decomposed into in-vehicle, waiting, walking, and transfer components and bounded by mode-specific transfer windows. On this representation a self-join path-enumeration engine constructs the feasible rail, air, and intermodal itineraries for every origin, destination, and departure, and a service-aware criticality measure ranks elements by the load of feasible least-cost itineraries they carry rather than by their position in the graph. Disruptions are applied through a single mechanism-driven operator that spans structural, institutional, and macro-scenario failures, so that topological and service-aware robustness are produced on the same network at the same scale and can be compared directly.

Third, building on the topological multilayer analysis of Ippolito and Cats (2024), the study extends continental robustness assessment from the structural layer to the service layer and uses it to analyse how the two readings of the same network diverge. It compares topological and service-aware criticality across the European rail and air networks, characterises the geographical distribution of the service-critical elements, and isolates the

institutional dimension of robustness by varying the cross-operator cooperation regime alongside the physical disruptions.

1.6. Thesis outline

The remainder of this thesis is organised as follows. Chapter 2 reviews the state of the art in long-distance transport networks, robustness analysis, and service-network representation, and identifies the methodological gaps addressed by the thesis. Chapter 3 introduces the multilayer network representation, the generalised travel-cost formulation, the path-finding algorithm for constructing passenger itineraries, and the three-type scenario taxonomy used to represent disruptions. Chapter 4 constructs the European rail and air multilayer network, reports its coverage and structural characteristics, and validates the path-enumeration engine against reference travel-time datasets. Chapter 5 reports the robustness results, evaluating the network under the structural, institutional, and macro-scenario disruptions of the taxonomy at both the topological and service layers. Chapter 6 answers the research questions, discusses the planning and policy implications of the findings, and outlines the limitations of the work and directions for future research. Figure 1.5 summarises this structure and maps each chapter to the research sub-question it addresses.

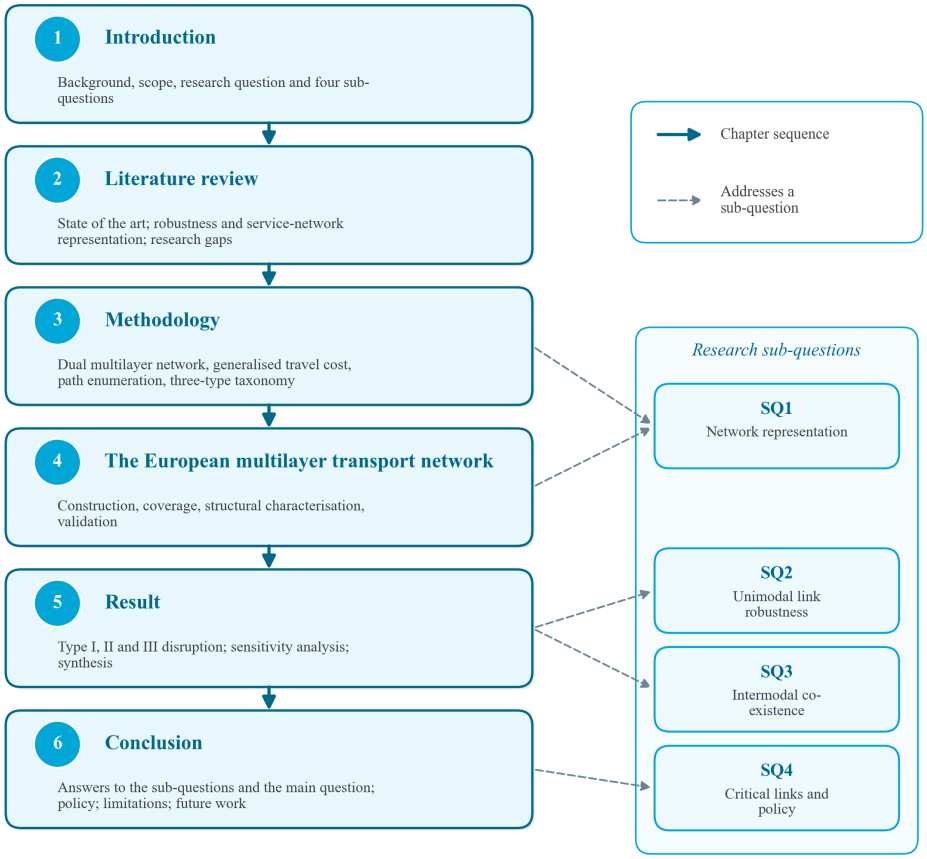

Figure 1.5: Thesis outline. The six chapters proceed from the research question to its answer, shown by the solid arrows, while the dashed links map each research sub-question to the chapters that address it.

2

Literature review

This chapter reviews the literature that situates and motivates the present study. Its organisation follows a single argumentative spine: topology determines whether a network stays connected, whereas service-aware analysis determines whether passengers retain usable itineraries. Every section points to one gap, namely that no service-weighted, mechanism-driven robustness analysis of the European rail and air networks yet exists at the continental scale. The chapter proceeds in three parts. Section 2.1 characterises long-distance travel in Europe and the structural fragmentation of its rail, air, and intermodal networks. Section 2.2 reviews the methodological literature on network robustness, tracing the transition from topological foundations to service-aware and mechanism-driven assessment. Section 2.3 synthesises the resulting research gaps and links each to a research sub-question formulated in Chapter 1. The generalised-travel-cost and transfer formulations that operationalise the service-aware perspective are part of this study's own modelling apparatus and are therefore developed in the methodology chapter rather than reviewed here.

2.1. Long-distance travel and its networks in Europe

Long-distance intercity travel is central to economic connectivity, climate policy, and territorial cohesion, yet it has historically received far less analytical attention than routine local mobility (Aultman-Hall & Ullman, 2020; van Goeverden et al., 2019). The imbalance is partly definitional and partly empirical: definitions of a long-distance journey may be spatial, temporal, or purpose-based (Malichová et al., 2022), and the most comprehensive EU-wide survey, the DATELINE project, dates back to 2001–2002 and predates the expansion of low-cost carriers, the construction of major high-speed lines, and EU enlargement (van Goeverden et al., 2019). The absence of detailed and reliable figures continues to hinder comparison and model validation at the European scale (Grolle et al., 2024). A spatial threshold of 100 km crow-fly distance is the most common convention (Aultman-Hall & Ullman, 2020; Frei et al., 2010; Van Goeverden, 2009), although such fixed cut-offs have

been criticised for aggregating very different trip types (Aultman-Hall & Ullman, 2020).[1]

For present purposes, the salient point is structural rather than statistical. The European long-distance market is dominated by three modes—car, rail, and air—whose balance shifts strongly with distance: the car and conventional rail prevail at shorter long-distance ranges, while air dominates the longer corridors where its cruising-speed advantage outweighs access and egress penalties (Grolle et al., 2024; Van Goeverden, 2009; van Goeverden et al., 2019). This distance dependence also governs the rail–air balance and the limited scope for modal shift. Brons et al. (2023) find that rail is faster than air on only 68 of 297 EU routes served by both modes, with rail competitiveness falling sharply beyond roughly 500 km, while Bruno et al. (2025) estimate that a ban on short-haul flights duplicated by sub-2:30 h rail would affect only 34 routes and 2.7% of intra-European air traffic. The same distance penalty constrains rail internally: cross-border services account for only about 5% of EU rail passenger-kilometres, and the number of cities connected by international rail declined between 2001 and 2019 (Cats, 2025; Eurostat, 2024).

The combination of sparse data, nationally oriented research, and the inherent complexity of cross-border, multi-operator, multimodal journeys means that the robustness of the European long-distance network at the continental scale has received almost no systematic attention. The remainder of this section reviews the structure and governance of that network; Section 2.2 then turns to the methods available for assessing its robustness.

2.1.1. Structural fragmentation of the European rail and air networks

A defining feature of the European long-distance system is that both of its surface-independent modes are fragmented, but for opposite reasons: rail along national infrastructural and institutional lines, air along market-structural lines within an otherwise unified regulatory and operational framework. This contrast is summarised in Table 2.1 and developed below, in each case closing on the implication for robustness.

Rail: infrastructural and institutional fragmentation

European rail evolved as a collection of nationally developed networks rather than a unified continental system (European Court of Auditors, 2018; UIC, 2018). Despite 175 years of development, interoperability remains constrained across some eight technical dimensions—gauge, signalling, electrification, operating speed, train length, loading gauge, homologation, and maintenance standards—only partially addressed by ERTMS and the Technical Specifications for Interoperability (Bairras & Aguas Ardaiz, 2025). Cats (2025) characterises the result as a "network-of-networks", and European Court of Auditors (2018) as an "ineffective patchwork" of national lines; tellingly, the number of cross-border trains hardly increased between 2001 and 2019 while the number of cities connected by international rail declined, a decline traced to the national perspective that planning authorities adopt when appraising cross-border links (Cats, 2025). Crucially, fragmentation extends beyond infrastructure: Raad voor de Leefomgeving en Infrastructuur (2020) decompose the international rail system into infrastructure, traffic

[1] Detailed market-size and modal-split figures for the European long-distance segment are reported in the introduction (Chapter 1); they are not repeated here.

management, transport services, and mobility services, and identify binding constraints—fragmented journey information, the inability to buy through-tickets across operators, and poor timetable coordination—at the passenger-facing layers. This multilayered fragmentation is directly relevant to robustness, because the usability of an alternative path after a disruption depends not only on whether infrastructure exists but on whether passengers can discover, book, and coordinate travel across operators.

A second structural property is spatial heterogeneity in network quality. Fraszczyk et al. (2016) identify three performance tiers across EU member states, with quality-of-service scores ranging from 7.1 (Switzerland) to 2.2 (Bulgaria) and investment strongly correlated with GDP per capita; Western and Central European networks are denser and more frequent, while Eastern and Southern systems are slower and underinvested. The robustness of the European rail network is therefore not spatially uniform, and a continental-scale assessment that ignores this heterogeneity risks averaging over fundamentally different vulnerability profiles. The high-speed layer compounds the picture: extensive within France, Spain, Germany, and Italy but poorly integrated across borders, it has been described as "poorly connected subnetworks" (Grolle et al., 2024), and its geographic reach remains an order of magnitude smaller than aviation's even at the hubs where the two modes are most closely integrated—in 2012, Paris CDG was linked to 234 cities by air but only 73 by HSR, and Amsterdam to 223 by air versus 12 by HSR (Dobruszkes et al., 2014).

Air: market-structural fragmentation under a unified framework

European aviation is fragmented in a fundamentally different way. Operating within a fully liberalised market since 1997, the air network is not a single graph but a superposition of strategically distinct sub-networks: hub-and-spoke legacy carriers that concentrate connectivity at a few nodes, and low-cost carriers that distribute it more evenly but at lower frequency per route (Adler et al., 2010; Dobruszkes, 2011). Since the mid-1990s, LCC expansion has grown the number of intra-European links substantially while leaving traffic highly uneven across them, so that the network has become denser in edges but increasingly heterogeneous in edge weight (Dobruszkes, 2011). Topologically, the network exhibits scale-free characteristics: Pien et al. (2015) model 1,254 airports and show that targeted removal of the highest-betweenness nodes fragments the giant component far faster than random failure, confirming the robust-yet-fragile property. From a robustness standpoint, the asymmetry between topological and volumetric structure implies that connectivity resilience may appear stronger than service-level resilience: removing a low-frequency LCC edge affects few passengers, whereas disruption at a major hub affects orders of magnitude more travellers while removing fewer edges.

A further asymmetry concerns operational coordination. Pellegrini and Rodriguez (2013) note that European airspace is managed through EUROCONTROL's Network Manager, providing continent-wide real-time flow management and rerouting, whereas rail remains dependent on nationally managed infrastructure with no equivalent body. This difference in coordination capacity is a structural factor affecting the realised robustness of each mode independently of its topology.

The intermodal relationship: competition, substitution, complementarity

The relationship between European air and rail is not one of simple substitution but a distance-dependent mix of competition, substitution, and emerging complementar-

ity (Aultman-Hall & Ullman, 2020). The best-documented interaction is the substitution of air by high-speed rail where rail travel time falls below roughly three hours: Dobruszkes (2011) shows near-complete air withdrawal on corridors such as Paris–Brussels once HSR time drops below about two hours, with air persisting or growing where it remains above three. At the continental scale, Dobruszkes et al. (2014) confirm across 161 city-pairs that HSR travel time is the single strongest predictor of air-service provision, with the sharpest impact between 2.0 and 2.5 hours, and caution that naive maximum-likelihood estimators overstate this effect by a factor of up to four. This substitution has been partially offset—and in aggregate outweighed—by LCC expansion, which Clewlow et al. (2014) find has added more medium- and long-haul traffic than HSR has removed at short range, so that HSR-induced shift on short corridors does not necessarily reduce system-level aviation emissions.

Beyond substitution, air and HSR are increasingly complementary at major hubs, where HSR acts as a feeder replacing short-haul spoke flights. F. Zhang et al. (2018) provide econometric evidence that on-site HSR integration raises total enplanement at primary hubs but not at secondary or regional airports, consistent with HSR extending a hub's catchment. From a network perspective this transforms competition into co-operation at specific nodes but introduces a new interdependency: disruption to HSR feeders can cascade into reduced air connectivity once the displaced short-haul flights are gone. Three robustness implications follow. First, on short-haul corridors where HSR has displaced air, modal redundancy has been removed, so a rail disruption causes absolute rather than redistributed connectivity loss. Second, at integrated hubs a new interdependency vulnerability emerges. Third, LCC-driven expansion has created medium- and long-haul air corridors with no rail alternative, leaving them dependent on the air network's structural robustness alone. These considerations motivate the integrated, multimodal treatment of robustness adopted in this study.

Table 2.1: Structural fragmentation of the European rail and air networks.

Dimension	Rail	Air
Infrastructure and technology		
Infrastructure control	National managers, vertical separation uneven	National airspace, pan-European coordination
Physical interoperability	Fragmented gauge, voltage, loading, homologation per country	Fully interoperable, any aircraft on any route
Signalling and control	20+ legacy systems, ERTMS ≈14% of core (2022)	Harmonised ATC under SES, single Network Manager
Market and operations		
Market access	4th Railway Package, open since Dec 2020, uneven uptake	Fully liberalised since 1997, free cabotage
Competition pattern	Incumbents dominant, open-access on dense trunks only	Hub-and-spoke legacy carriers vs. continental LCCs
Cross-border share	≈7% of pkm, connected cities declined 2001 to 2019	Dominant cross-border mode
Crew and labour	National licences, language on cab, per-country homologation	Common EASA licence, English standard, portable crews
Customer-facing integration		
Ticketing and information	No mandatory through-tickets, fragmented platforms	IATA interlining, global distribution systems
Disruption accountability	Diffuse, liability unclear across borders	Centralised flow control, uniform passenger rights
Governance		
Network coordinator	None, ERA regulates but does not operate	EUROCONTROL Network Manager
Policy framework	Single European Railway Area, partial	Single European Sky, active since 2004

Synthesised from Bairras and Aguas Ardaiz (2025), Beria et al. (2023), Cats (2025), Pellegrini and Rodriguez (2013), Tomeš and Pařil (2026), and Witlox et al. (2022) and (European Commission, 2025).

2.2. Robustness analysis of transport networks

2.2.1. Definitions of robustness

Robustness carries multiple, context-dependent meanings. In network science it is broadest: the capacity of a network to maintain global connectivity when nodes or edges are removed through failure or attack (Albert et al., 2000; Barabási, 2016). In railway operations it is narrower, denoting the capability to absorb everyday delays through slack and buffers (Bešinović, 2020); in air transport it is often framed as the ability to sustain traffic flows under capacity degradation (Pien et al., 2015) or efficiency weighted by frequency or passenger volume (Zhou et al., 2019). A growing strand defines robustness functionally, as the degree to which a network maintains its level of service—measured through passenger-relevant metrics such as generalised travel cost, accessibility, or welfare—under disruption (Faturechi & Miller-Hooks, 2015; Mattsson & Jenelius, 2015). This distinction

is consequential: a purely topological analysis can establish whether alternative paths exist but not how costly they are, whereas a functional, service-aware analysis quantifies the passenger-experienced consequences. The latter has been increasingly advocated for large-scale assessment (Bešinović, 2020; Jiao et al., 2020), yet, as the following subsections show, most studies remain topological and none applies service-aware metrics to the European networks at continental scale.

2.2.2. Topological approaches and their limits

Complex-network theory offers a common level of abstraction for placing modes and regions in a single analytical frame: nodes denote stations, airports, or cities, and edges denote physical or service links, enabling indicators comparable across networks of different scale and geography (Barabási, 2016; von Ferber et al., 2009). In a review of 154 papers, M. Zhang et al. (2022) find that robustness studies constitute about one-third of complex-network applications in transport. A topological robustness analysis removes nodes or edges according to a strategy and tracks a performance metric as the network degrades. The metrics actually used in this study are the size of the giant connected component S and the global efficiency of Latora and Marchiori (2001),

$$E(G) = \frac{1}{n(n-1)} \sum_{i \neq j} \frac{1}{d(i,j)}, \tag{2.1}$$

where $1/d(i,j) = 0$ when no path exists, so $E(G)$ degrades gracefully as the network fragments and, unlike average path length, does not require connectivity; the relative efficiency loss $\Delta E/E_0$ is the standard robustness indicator in this literature (Ippolito & Cats, 2024; Pien et al., 2015). Targeted strategies rank elements by centrality—degree, betweenness, or closeness—which also identifies the elements whose protection matters most (Mattsson & Jenelius, 2015; M. Zhang et al., 2022).

Robustness behaviour is conditioned by the degree distribution. Scale-free networks, whose degree distribution follows a power law, are robust to random failure but fragile to targeted hub removal—the "robust-yet-fragile" property established by Albert et al. (2000) and among the most cited results in network science (Barabási, 2016). Transport networks generally exhibit this behaviour, though their degree distributions are often truncated by physical and operational constraints (von Ferber et al., 2009; M. Zhang et al., 2022). Two canonical disruption strategies dominate: random failure, a lower bound on vulnerability, and targeted attack, an upper bound; robustness is commonly summarised by the area under the degradation curve, a single index $R \in [0,1]$ (Albert et al., 2000; Ippolito & Cats, 2024; Pien et al., 2015).

These tools have proven productive at continental scale. For European rail, Calzada-Infante et al. (2020) characterise the international railway network (412 cities) and find a power-law degree distribution with small-world properties; critically, once weighted by service frequency, the highest-degree node (Paris) diverges from the highest-strength node (Frankfurt)—an early empirical signal that topology alone may misidentify critical elements. Ippolito and Cats (2024) analyse European rail and air jointly across 124 metropolitan areas using binary adjacency, finding rail substantially less robust than air under targeted attack owing to its corridor-based topology, but explicitly acknowledge that their binary representation omits temporal and operational attributes and call for time-dependent

service graphs "where each link corresponds to a specific train departure".

That call reflects three documented limits of topological analysis. First, importance rankings change once service information is incorporated, as the Paris–Frankfurt divergence and the analogous frequency-weighted reranking of airports by Zhou et al. (2019) both show; protection priorities derived from topology alone may be misallocated (Bešinović, 2020). Second, static metrics fail precisely during disruption: Bešinović (2020) reports that static and dynamic measures agree under normal conditions but diverge once interdependencies and passenger flows are activated by a disruption. Third, binary connectivity ignores service-quality degradation—a maintenance-affected corridor still carries trains at reduced frequency, a state that edge removal cannot represent, so topological analysis either overstates damage (removing a degraded-but-usable edge) or understates it (retaining an edge whose service has fallen below usability) (Cats et al., 2017; Mattsson & Jenelius, 2015). Bešinović (2020) accordingly distinguishes *topological* metrics, which assess static structure, from *system-based* metrics, which represent demand, supply, and the response to and recovery from disruption, and which include generalised travel cost, rerouting rates, and the number of passengers unable to reach their destination within an acceptable time.

2.2.3. Service-aware robustness

A growing body of work embeds service information—timetables, generalised travel cost, rerouting behaviour, and operator structure—into robustness assessment. The simplest step replaces binary adjacency with service-weighted edges: Zhou et al. (2019) weight air-network edges by route count, frequency, and capacity and show that weighting alters both efficiency and the identity of critical airports, while Jiao et al. (2020) introduce a Weighted Network Efficiency for high-speed rail combining frequency and travel time, finding that 10% betweenness-targeted removal causes more than 10% efficiency loss whereas 80% random removal is needed for the same effect, and that disruption impact varies strongly with time of day—something topological metrics cannot capture.

A more fundamental extension concerns the nature of the disruption. Most real disruptions cause *partial* capacity reduction rather than complete failure (Cats & Jenelius, 2018). Cats and Jenelius (2018) reduce line frequencies in 25% increments on Stockholm's rapid-transit network and find that societal cost rises *convexly*: across metro lines a 75% reduction causes on average 8.5 times the damage of a 25% reduction, not three times, driven by non-linear waiting time, denied boarding, and cascading reroute congestion—and ridership alone does not predict criticality, confirming that vulnerability is a network rather than a local property. The companion work of Cats et al. (2017) formalises this into a *link criticality* $c_e = \int_0^1 [y_e(x) - y_e(0)]\,\mathrm{d}x$ and a *degrading rapidity* r_e capturing the shape of the normalised degradation curve, with rapidity below 0.5 indicating convex degradation; for the Amsterdam network the mean rapidity of 0.26 confirms that convexity is the norm.

The most comprehensive departures are time-varying and multimodal. Bellocchi et al. (2021) propose a dynamical efficiency that tracks performance at fine temporal resolution and, via realistic intermodal paths, find it correlates only weakly with link speed and barely at all with betweenness, confirming it captures distinct information. Xu and Chopra (2023) add a three-indicator resilience framework and find, counter-intuitively, that integrating a *vulnerable* subsystem can improve overall robustness provided its intermodal connec-

tions are strong. Wang et al. (2025) weight a multimodal network by both generalised travel cost and flow, with user-equilibrium redistribution, and show that efficiency deteriorates non-linearly over time, reinforcing the convexity result in a flow-dependent setting. Across these studies four findings recur: weighting changes which elements are critical; vulnerability is convex; resilience involves recovery shape and speed, not just loss magnitude; and multimodal interconnection is a resource only when transfers are practically usable. Despite this trajectory, no study combines scheduled service times, GTC decomposed into in-vehicle, waiting, and transfer components, and operator-level structure with progressive robustness analysis at the European continental scale, for rail alone let alone the integrated air–rail network.

2.2.4. Mechanism-driven disruption scenarios

The two canonical strategies are useful bounds but correspond to no recognisable disruption mechanism: random failure assumes spatially uniform probability, and targeted attack assumes perfect adversarial knowledge (Albert et al., 2000; M. Zhang et al., 2022). Real disruptions are structured—floods affect contiguous corridors, strikes affect one operator's services, border closures affect specific country pairs, maintenance affects defined segments for defined periods (Mattsson & Jenelius, 2015)—so abstract attack findings ("the network fragments when the top 10% of hubs are removed") do not map onto actionable scenarios. A small literature has begun to model identifiable mechanisms: Bešinović (2020) reviews data-driven studies using historical hazard records, and Jiao et al. (2020) model a COVID-inspired cascade on the Chinese HSR, though with stylised adjacent-node propagation. At the European scale, Ippolito and Cats (2024) again use only random and targeted strategies on a binary graph; no European-scale study has implemented operator withdrawal, corridor closure, or border closure on a service-weighted network.

These structured mechanisms are not hypothetical, and the European system has seen a clear instance of each in recent years. A spatially contiguous hazard removes every service along an affected stretch of infrastructure at once, as the July 2021 floods did across the Rhine and Ahr valleys (Koks et al., 2022; Mohr et al., 2023) and a volcanic-ash cloud did across much of European airspace in 2010 (Eurocontrol, 2010; Ulfarsson & Unger, 2011). An institutional disruption removes a single operator's services while leaving the track and the runway intact, a recurrent strain on a network whose cross-border connectivity stays fragmented across national incumbents under still incomplete interoperability (Bairras & Aguas Ardaiz, 2025; European Union Agency for Railways, 2025). The loss of a major station or airport removes both its direct services and the transfers that stitch separate operators together, while the redrawing of the trans-European corridors after 2022 shows how fast a cross-border connection can be reconfigured by events outside the transport sector (European Commission, 2022). What these mechanisms share is that none can be reproduced by removing a random or a maximally central node, since each is defined by operator, geography, or timetable (Mattsson & Jenelius, 2015), and so calls for exactly the service-level information that a topological model discards. Toet et al. (2026) sharpen why this matters, showing that routing which exists on paper in European long-distance networks can stay out of a traveller's reach when booking and disruption management are not coordinated across operators, which leaves open whether the redundancy a topological graph displays is redundancy a traveller can actually use.

2.3. Research gaps

The preceding review reveals four interlocking gaps, spanning a *methodological* axis (how robustness is measured) and a *scope* axis (what is measured). Each is stated below in one paragraph and linked to the sub-question it motivates; Table 2.2 summarises the correspondence.

Gap 1 : Representation. Continental robustness studies of European long-distance networks represent edges as binary indicators of service existence, omitting timetable, stop sequence, and transfer constraints (Ippolito & Cats, 2024; Pien et al., 2015). Consequently the shortest path is defined by leg count rather than travel time or service quality, and Ippolito and Cats (2024) explicitly call for time-dependent graphs. No timetable-based, time-dependent representation yet exists for the European continental-scale air and rail networks. This motivates Sub-question 1.

Gap 2 : Metric. Even studies that move beyond binary topology evaluate robustness with graph-theoretic indicators—degree, betweenness, hop-count efficiency—that do not reflect passenger-experienced service quality (Ippolito & Cats, 2024; Pien et al., 2015). The generalised-travel-cost framework decomposes a journey into in-vehicle, access, waiting, and transfer components with differentiated weights, and transfers in particular are a major deterrent to long-distance rail (de Pindray d'Ambelle, 2024); yet no robustness study has measured degradation in GTC terms, which would capture the loss of direct connections and forced reroutes rather than mere topological change. This motivates Sub-question 2.

Gap 3 : Scale. No service-weighted robustness analysis of the European rail and air networks exists at continental scale: Pien et al. (2015) and Ippolito and Cats (2024) are topological, while the timetable-based work of Jiao et al. (2020) is confined to Chinese HSR. The gap is compounded by a pervasive high-speed-rail bias that neglects the conventional network connecting intermediate and peripheral cities (Clewlow et al., 2014; Dobruszkes, 2011), and by the unaddressed geographic heterogeneity of European rail quality (Fraszczyk et al., 2016), both of which a continental analysis must accommodate to avoid averaging over distinct vulnerability profiles. This motivates Sub-questions 2 and 3.

Gap 4 : Intermodal redundancy. The air–rail relationship combines substitution, complementarity, and growing interdependence (Clewlow et al., 2014; Dobruszkes et al., 2014; F. Zhang et al., 2018), yet whether rail–air co-existence supplies *usable* rather than merely formal redundancy has not been tested at the service level. Whether the timetable, frequency, and transfers of one mode can absorb disrupted demand from the other—and how this is conditioned by inter-operator coordination (O'Sullivan & Patel, 2004) and by the asymmetry in operational coordination between the modes (Pellegrini & Rodriguez, 2013)—remains unknown. This motivates Sub-questions 3 and 4.

Table 2.2: Summary of identified research gaps and their correspondence to research sub-questions

Gap	Description	Addressed by
1	Binary network representations omit service-layer dynamics (timetable, stop sequences, transfer constraints)	Sub-question 1: multilayer graph representation capturing infrastructure and service layers
2	Robustness metrics do not reflect passenger-experienced service quality (no GTC-based assessment)	Sub-question 2: robustness analysis incorporating service-level information
3	No European-scale service-weighted robustness analysis; HSR bias and geographic heterogeneity unaddressed	Sub-question 2 and 3: unimodal and intermodal robustness at continental scale
4	Intermodal robustness and modal redundancy untested at service level; coordination asymmetry unquantified	Sub-question 3 and 4: intermodal integration and planning implications

3

Methodology

This chapter describes the methodological framework for assessing the robustness of the European MLTNs, corresponding to the complex network representation of the real-world European intercity rail and air network. Section 3.1 defines the network representation and its spatial resolution. Section 3.2 models travel time on the network, developing the generalised travel cost in Section 3.2.4 and the path enumeration in Section 3.2.5. Section 3.3 presents the robustness assessment metrics and indicators. Section 3.4 establishes the scenario taxonomy that structures the experiments reported in Chapter 5.

3.1. Network representation

As established in Section 1.3.2, this study distinguishes between two complementary multilayer transport-network representations. The structural MLTN captures the topological organisation of the European rail and air system, with rail and air modelled as separate but interdependent modal layers. The structural MLTN resembles the pure-topological study as in Ippolito and Cats (2024). The functional MLTN builds on the same physical system but adds supply-side information, namely scheduled timetables, operating calendars, and the ordered station or airport sequences served by each scheduled service, from which the generalised travel cost is derived. Both representations share the same physical entities, but they differ in the amount and type of information carried by their edges. The following subsections describe how the two representations are constructed, covering the node set, edge set, inter-layer coupling, transfer rules, timetable-based feasibility constraints, and disruption mechanisms.

3.1.1. Spatial resolution and urban representation

Traffic modelling is, in essence, a simplified representation of a real-world transport system, and transport models are therefore inherently problem-specific and viewpoint-specific (Ortúzar & Willumsen, 2011). Under the principle of variable-resolution modelling, the network should be represented at a level of detail fine enough to retain the mechanisms that are relevant to the analytical objective. Since this study evaluates ro-

bustness through the availability of alternative paths between origin–destination pairs, the network representation must preserve the route choices, transfer opportunities, and station-level connections that passengers may use when their primary path is disrupted. Aggregating stations and airports directly into city-level nodes would merge distinct transfer points and terminal locations into a single abstract node, thereby masking the station-specific vulnerabilities that shape service-layer robustness. For this reason, both structural MLTN and functional MLTN networks are defined at the station and airport level, allowing intra-urban heterogeneity in connectivity and transfer opportunities to be explicitly captured.

Although the network is represented at the station and airport level, a spatial clustering rule is required to distinguish local access and transfer movements within the same urban area from long-distance movements between different urban areas. The Functional Urban Area (FUA) provides an appropriate spatial unit for this purpose because it captures the city together with its functionally connected commuting zone rather than relying only on administrative boundaries. Following the Eurostat–OECD definition (Dijkstra et al., 2019), an FUA consists of a city and its commuting zone. The city is derived from a high-density urban centre, which is a cluster of contiguous 1 km² grid cells with at least 1,500 inhabitants per km² and a total population of at least 50,000. It includes the local administrative units in which at least 50% of residents live within that centre, while the commuting zone comprises surrounding contiguous units from which at least 15% of employed residents commute to the city (Dijkstra et al., 2019).

This choice is particularly appropriate because passenger-side demand data are outside the scope of the study. In the absence of observed demand flows, the FUA provides a functional proxy for the urban catchment of each station or airport, since it is defined not only by population concentration but also by commuting relations between the city and its surrounding local units. This makes the FUA more suitable than purely administrative regional units, such as NUTS regions (European Parliament and Council of the European Union, 2003), for defining origin and destination zones in an intermodal passenger network. Whereas NUTS regions are primarily hierarchical statistical regions based on administrative territorial units, FUAs approximate the labour-market catchment of an urban area. This functional property is important here because a railway station or airport may serve a population distributed across the surrounding commuting zone, rather than only the population located within the same administrative boundary.

A further justification for using FUAs is the cross-national scope of the study. As stated in Section 1.3, the study area includes several countries outside the European Union, where administrative city boundaries may be defined differently and are therefore not directly comparable. FUAs, by contrast, are based on a standardised functional definition that reflects the relationship between an urban centre and its surrounding commuting catchment. They therefore provide a consistent basis for comparing robustness metrics across European urban areas with different administrative systems, and for distinguishing local intra-urban movements from intercity long-distance relations within a common spatial framework. Additionally, defining inter-city services by whether they connect different FUAs avoids imposing a rigid distance threshold, such as classifying all services above 100 km as long-distance. This is particularly important in a cross-national study, where country size, urban density, and settlement patterns vary substantially.

Lastly, the FUA is often the proper spatial scale at which rail and air modes naturally coexist. Airports are rarely located in administrative city centres, but they are almost always within the Functional Urban Area. As Kristoffersson and Berglund (2022) observe, airports are typically located further from city centres than main train stations, and connection trips account for 30–50% of total trip cost for air travel compared to less than 10% for rail. To analyse multi-modal robustness, one must use a spatial scale that encompasses both the airport and the major rail stations and the FUA is the standard unit that captures this interaction.

In sum, this study adopts the FUA as the primary unit of OD aggregation to ensure cross-nation statistical consistency, to capture the full catchment area of international transport hubs, and to encompass both rail and air terminals within a single spatial unit. Figure 3.1 summarises this spatial framework, with the European-scale panel showing the continental distribution of rail and air terminals across the 387 FUAs of the study area, and the Paris panel illustrating the intra-FUA terminal arrangement.

3.1.2. Multilayer representation of the transport network

As introduced in Section 1.3.2, this study develops two complementary multilayer transport network representations, namely a structural MLTN and a functional MLTN. The structural MLTN distinguishes between the rail and air layers and abstracts the physical infrastructure and operational connection into simplified network edges. Together, these structural edges define the topological potential for unimodal rail services, unimodal air services, and intermodal rail–air connections. The functional MLTN, by contrast, uses the structural MLTN as its topological substrate and extends it with service-based and temporal information on supply layer. This includes service schedule, operators, departure and arrival times, transfer constraints, and feasible departure windows. The functional MLTN therefore serves as the basis for the path-enumeration procedure, through which feasible origin–destination itineraries are constructed and evaluated in terms of generalised travel cost.

These two representations are kept apart rather than merged into one graph, as a deliberate separation of concerns. The structural MLTN records only whether a connection exists, on which mode, and where the two modes interchange, and it is the level on which the robustness experiments of Chapter 5 remove edges, since structural criticality is best measured on the bare connectivity skeleton. The functional MLTN adds the supply detail of scheduled services with their operators, timetables, and calendars, from which the generalised travel cost is derived and on which the path enumeration and the operator-withdrawal experiments are built. Separating the two lets a disruption be aimed at the structural or the supply level independently, while their common node set keeps the two sets of results comparable.

3.1.3. Mode-layered structural MLTN

Following the taxonomy of Šfiligoj et al. (2026), the structural mode-layered representation is an **interdependent multilayer network** in which each layer encodes the mode-specific topological connections between facilities. In the rail layer, edges represent direct connections between consecutive railway stations, thereby simplifying the underlying track infrastructure. In the air layer, edges represent direct airport-pair connections, abstracting

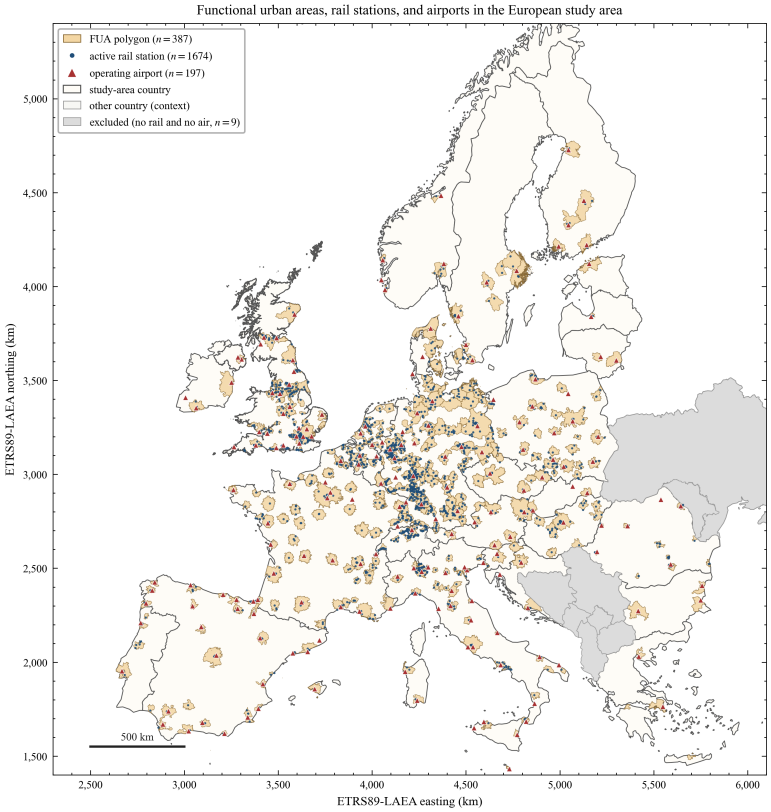

(a) European-scale FUA node representation, showing the 387 FUA polygons (beige), 1,674 active rail stations (blue dots), and 197 operating airports (red triangles) retained in the analysis. The countries with neither active rail nor air service are excluded (grey-shaded).

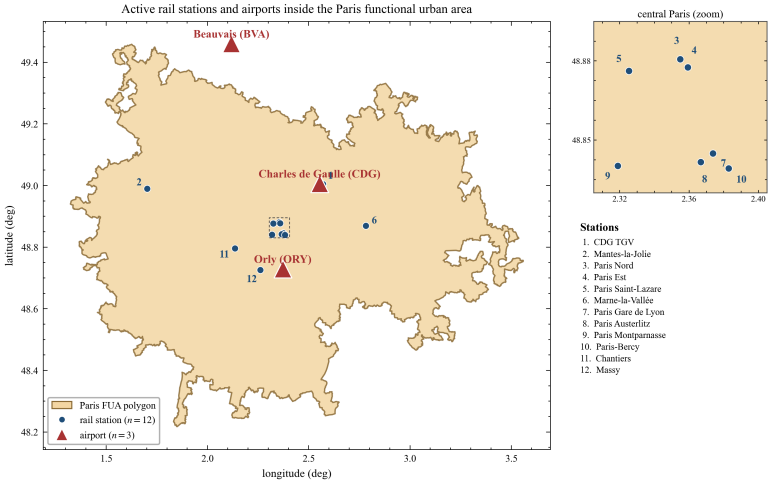

(b) Paris FUA as a representative case, with the twelve intercity rail stations (blue dots) and three commercial airports (Charles-de-Gaulle, Orly, Beauvais; red triangles) enclosed by the FUA polygon. The inset enlarges the central-Paris cluster of terminal stations that are indistinguishable at the European scale.

Figure 3.1: Functional Urban Area-based spatial framework of the European multilayer transport network. The European-scale panel (a) shows how stations and airports are assigned to FUAs across the study area, while the Paris panel (b) illustrates the spatial distribution of rail and air terminals within a single FUA.

from the detailed airspace trajectories through which flights may be operated. The focus is structural and aligns with the classical topological representation of Ippolito and Cats (2024). Each active edge in the structural MLTN is assigned a unit weight, representing the existence of an operational connection between two stations or airports, while generalised travel cost could be introduced in the functional MLTN described in Section 3.1.4. The structural MLTN is formally specified as

$$G^{\mathrm{s}} = \big(V, E^{\mathrm{s}}, L^{\mathrm{s}}, \mathscr{F}, \mathbf{1}\big), \qquad E^{\mathrm{s}} = E^{r}_{\mathrm{intra}} \cup E^{a}_{\mathrm{intra}} \cup E_{\mathrm{inter}}, \tag{3.1}$$

where $L^{\mathrm{s}} = \{r, a\}$ indexes the rail and air layers, $V = V^r \sqcup V^a$ is the disjoint union of $|V^r| = 1{,}674$ intercity rail stations and $|V^a| = 197$ airports which have active services traversing through. The superscript on L^{s} disambiguates the structural mode-layer index from the functional aspect-layer index $L^{\mathrm{f}} = \{\mathrm{topo}, \mathrm{sup}\}$ introduced in Section 3.1.4. $\mathscr{F}$ is the set of Functional Urban Area polygons defined in Section 3.1.1 and shown in the Figure 3.1a. Each node $v \in V$ is mapped to its containing FUA through

$$\phi : V \longrightarrow \mathscr{F} \cup \{\varnothing\}, \qquad \phi(v) = \begin{cases} f & \text{if } v \text{ belongs to FUA } f \in \mathscr{F}, \\ \varnothing & \text{otherwise.} \end{cases} \tag{3.2}$$

The mapping ϕ provides the spatial basis for the subsequent edge-construction rules by identifying whether two stations or airports belong to the same FUA. The component edge sets E^{r}_{intra}, E^{a}_{intra} and E_{inter} are defined by equations 3.5–3.7. Lastly, when the structural graph is used only for topological analysis, the weighting function $\mathbf{1} : E^{\mathrm{s}} \to \{1\}$ assigns the same unit weight to all edges. The two mode layers are described in turn below, followed by the interdependent coupling rule and the supra-adjacency assembly and the Figure 3.2's visualization.

Rail layer of the structural MLTN

The rail layer is the unweighted simple graph $G^r = (V^r, E^{r}_{\mathrm{intra}})$ whose edges record the direct connection between two consecutive stations served by an operating rail service. Following the spatial-scope distinction of Section 3.1.1, the set of services are split into two classes by the FUA membership of their stops:

$$\mathscr{S}^r = \mathscr{S}^r_{\mathrm{loc}} \cup \mathscr{S}^r_{\mathrm{ld}}, \tag{3.3}$$

where $\mathscr{S}^r_{\mathrm{loc}}$ collects the intra-FUA local-public-transport services that link station pairs within the same urban catchment. These services may include commuter rail, metro, tram, bus, or walking links and $\mathscr{S}^r_{\mathrm{ld}}$ collects the long-distance rail services loaded from GTFS feeds. Each long-distance service $s \in \mathscr{S}^r_{\mathrm{ld}}$ carries an ordered stop sequence $\sigma(s) = (\sigma_1(s), \ldots, \sigma_{n_s}(s))$, on which the consecutivity relation

$$v_i \sim_s v_j \iff \exists k \in \{1, \ldots, n_s - 1\} : \{\sigma_k(s), \sigma_{k+1}(s)\} = \{v_i, v_j\} \tag{3.4}$$

is defined. Intra-urban services $s \in \mathscr{S}^r_{\mathrm{loc}}$, by contrast, are abstracted as single terminal-to-terminal links and are characterised only by their endpoint pair $\eta(s) = \{v^r_i, v^r_j\}$ without an intermediate stop sequence. The two service classes together instantiate an undirected rail edge through

$$\left(v_i^r, v_j^r\right) \in E^r_{\text{intra}} \iff \underbrace{\left[\phi(v_i^r) = \phi(v_j^r) \neq \emptyset \;\wedge\; \exists s \in \mathcal{S}^r_{\text{loc}} \,:\, v_i^r \sim_s v_j^r\right]}_{\text{intra-FUA rule}} \;\vee\; \underbrace{\left[\exists s \in \mathcal{S}^r_{\text{ld}} \,:\, v_i^r \sim_s v_j^r\right]}_{\text{long-distance rule}}. \tag{3.5}$$

The two disjuncts produce the labelled decomposition

$$E^r_{\text{intra}} = E^r_{\text{intra,loc}} \cup E^r_{\text{intra,ld}}, \tag{3.6}$$

in which each edge carries a flag indicating which disjunct produced it. An $E^r_{\text{intra,ld}}$ edge therefore represents two things at once. It is both an operational connection, in the sense that at least one service in $\mathcal{S}^r$ stops at v_i^r and at v_j^r in immediate succession and a simplified representation of the physical infrastructure those services share, in the sense that the different train services operating over the same track between two stations collapse to a single edge, regardless of the actual track the services use. Edges in $E^r_{\text{intra,loc}}$ follow analogous logic but are abstracted as direct terminal-to-terminal links as they carry no intermediate-stop sequence and no station ordering, regardless of how many stops the underlying urban service traverses. The spatial and temporal effort of those intermediate stops is absorbed into the connection-time attribute $t^{\text{connect}}_{\text{intra-urban}}$ of Section 3.2.2. The same abstraction applies to the ground-access edges E_{inter} introduced in the inter-layer coupling.

Air layer of the structural MLTN

The air layer is also a unweighted simple graph $G^a = (V^a, E^a_{\text{intra}})$, constructed by the same rule with $\mathcal{S}^a = \mathcal{S}^a_{\text{loc}} \cup \mathcal{S}^a_{\text{ld}}$. Here $\mathcal{S}^a_{\text{loc}}$ collects the intra-FUA airport-to-airport ground-transfer services that link multiple airports within the same urban catchment, while $\mathcal{S}^a_{\text{ld}}$ denotes the set of scheduled commercial air services obtained from airline timetable data. Every flight in $\mathcal{S}^a_{\text{ld}}$ is a non-stop service with $n_s = 2$, so the consecutivity relation $\sim_s$ collapses to a single edge between origin and destination airport; the same convention applies to single-hop services in $\mathcal{S}^a_{\text{loc}}$. Equation 3.5 specialises to $\ell = a$ and yields the same labelled decomposition $E^a_{\text{intra}} = E^a_{\text{intra,loc}} \cup E^a_{\text{intra,ld}}$ as in the rail layer. Multiple airlines flying the same long-distance airport pair collapse to a single edge in $E^a_{\text{intra,ld}}$ and the per-airline strands are recovered later via the operator overlay $\Pi(e)$ of equation 3.13. In other words, each $E^a_{\text{intra,ld}}$ edge also aggregates into a single link the airspace trajectories flown by the individual services between the two airports. The air-layer adjacency $A^a \in \{0,1\}^{|V^a|\times|V^a|}$ is the corresponding 0/1 matrix.

Interdependent coupling

The two layers are coupled by edges between a rail station and an airport that share an FUA and are linked by a scheduled ground-access service or, in the simplest case, by a walkable distance within the same terminal complex. Let $\mathcal{S}_{\text{trans}}$ denote the set of such ground-access services (dedicated airport rail link, branded shuttle, suburban-rail stop inside the terminal complex, walking). An inter-layer edge satisfies

$$\left(v_i^r, v_j^a\right) \in E_{\text{inter}} \iff \phi(v_i^r) = \phi(v_j^a) \neq \emptyset \;\wedge\; \exists s \in \mathcal{S}_{\text{trans}} \,:\, v_i^r \sim_s v_j^a, \tag{3.7}$$

mirroring equation 3.5 in form but operating across layers. Every ground-access service in $\mathcal{S}_{\text{trans}}$ is a terminal-to-terminal link with $n_s = 2$, so each s produces a single edge between its rail and air endpoints. Inter-layer edges are treated as undirected in the structural representation and the directional walking-and-transit time recorded by the urban routing engine, which is generally asymmetric between any pair of co-located rail station and airport, is reintroduced in the functional MLTN of Section 3.1.4 through the GTC weighting of Section 3.2.2. The coupling block of the supra-adjacency below is $C^{ra} \in \{0,1\}^{|V^r|\times|V^a|}$, with $C^{ra}_{ij} = 1$ if the rule of equation 3.7 is satisfied.

Supra-adjacency form

The full edge set of the structural MLTN is the disjoint union

$$E^{\text{s}} = E^{r}_{\text{intra}} \cup E^{a}_{\text{intra}} \cup E_{\text{inter}}, \tag{3.8}$$

and the structural MLTN admits the block representation (Šfiligoj et al., 2026, eq. (1))

$$\mathbf{A}^{\text{s}} = \begin{bmatrix} A^r & C^{ra} \\ C^{ar} & A^a \end{bmatrix}, \tag{3.9}$$

in which the diagonal blocks A^r and A^a encode the intra-layer rule of equation 3.5 at $\ell = r$ and $\ell = a$ respectively, and the off-diagonal blocks $C^{ra} = (C^{ar})^\top$ encode the inter-layer rule of equation 3.7. Equation 3.9 is the operand for the topological robustness experiments of Chapter 5 and is the structural counterpart of the multiplex supra-adjacency $\mathbf{A}^{\text{f}}$ of equation 3.16.

3.1.4. Supply-layered functional MLTN

Whereas the structural MLTN in Section 3.1.3 records the *existence* of direct connections between transport terminals, the functional MLTN characterises the scheduled services that realise these connections, including the operator, departure window, and service duration. Rather than consolidating both representations into a single graph, the two graphs are maintained as complementary objects. This separation keeps the representation tractable, avoiding the analytical and notational burden of a heavily interwoven four-layer structure. Modal distinctions remain recoverable through the structural MLTN, and disruption experiments can be targeted at the topological or supply layer independently. Following the taxonomy of Šfiligoj et al. (2026, §2.2), this is most naturally a **supply-layered multiplex MLTN**, because its two layers share the same node set and are coupled by the identity mapping between node replicas, with each node joined only to its own counterpart on the other layer. The functional MLTN is the multiplex denoted as

$$G^{\text{f}} = \left(V, E^{\text{f}}, L^{\text{f}}, \tau, w^{\text{f}}\right), \qquad L^{\text{f}} = \{\text{topo}, \text{sup}\}, \tag{3.10}$$

where $V = V^r \sqcup V^a$ is the node set inherited verbatim from the structural MLTN, $\tau : V \to \{r, a\}$ records the mode attribute. The two layers are described in turn below.

Topological layer

This layer is a flattened rendering of the structural MLTN with mode demoted from layer index to node attribute. Because the mode-layered supra-adjacency $\mathbf{A}^{\text{s}}$ of equation 3.9 is

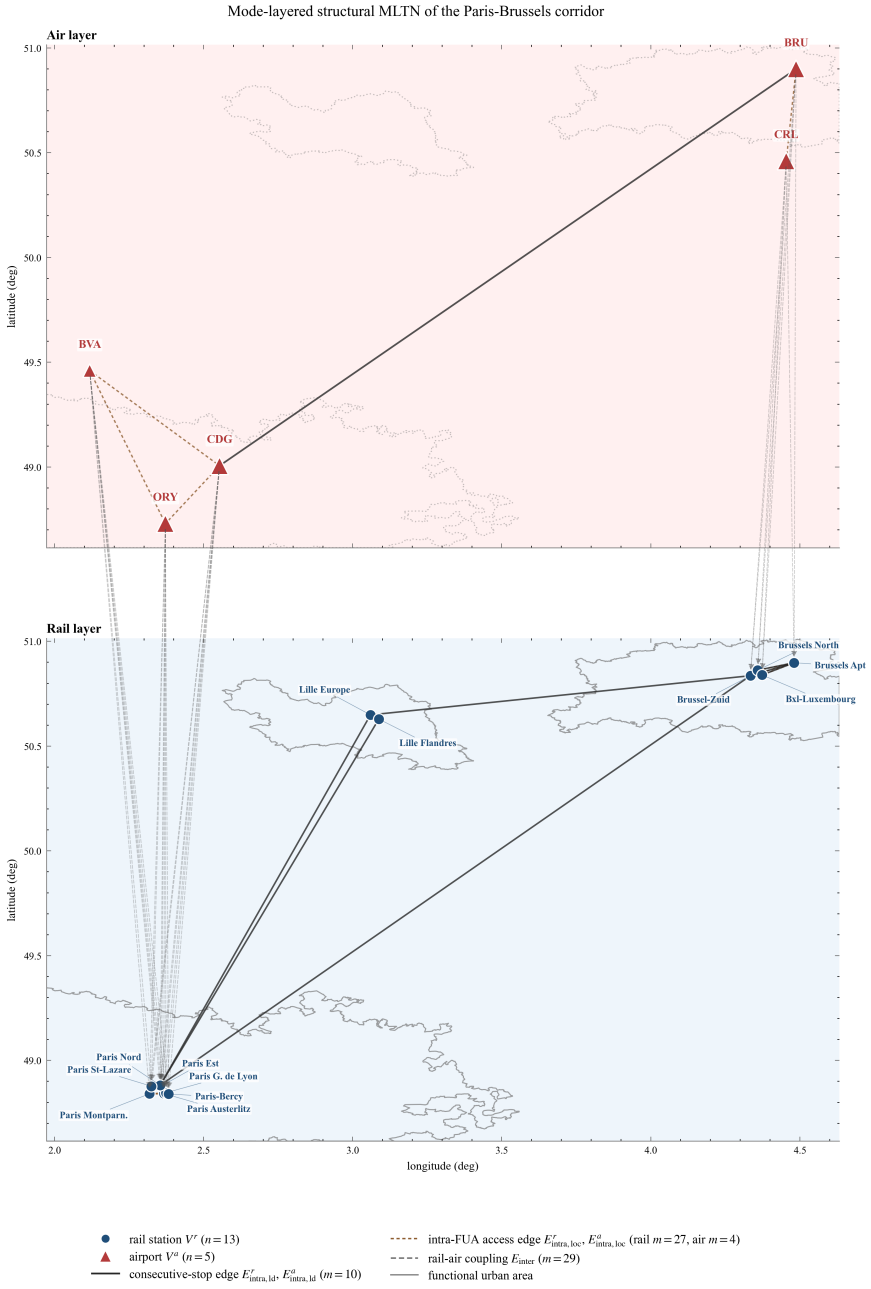

Figure 3.2: Mode-layered structural MLTN of the Paris-Brussels corridor, spanning the Paris, Lille, and Brussels/Brussel functional urban areas. The air layer (top pane) contains five airports (ORY, CDG, BVA, BRU, CRL) and the rail layer (bottom pane) contains thirteen intercity stations on the Eurostar high-speed backbone between Paris Nord, Lille Europe, and Brussel-Zuid. Solid black edges encode the consecutive-stop adjacencies of equation 3.5 ($m = 9$ on the rail layer and $m = 1$ on the air layer). Dashed brown edges aggregate intra-FUA connections within the urban access subgraph ($m = 27$ station-to-station and $m = 4$ airport-to-airport). Dotted grey inter-layer edges represent the $m = 29$ rail-air footpath connections of equation 3.7, which bind the two infrastructure layers at airport-rail interchanges such as CDG TGV, Brussels Airport rail, and Schiphol rail.

already block-square over the joint node set $V = V^r \sqcup V^a$, demoting the layer index to a node attribute leaves the matrix unchanged in form. The topological adjacency is therefore the structural supra-adjacency itself.

$$A^{\text{topo}} = \begin{bmatrix} A^r & C^{ra} \\ C^{ar} & A^a \end{bmatrix} = \mathbf{A}^{\text{s}}, \qquad E^{\text{topo}} = E^r_{\text{intra}} \cup E^a_{\text{intra}} \cup E_{\text{inter}}, \tag{3.11}$$

All edges are assigned unit weight, so the topological layer preserves the binary connectivity of the structural MLTN and the only modification is representational, with the two modal layers collapsed onto a common node set indexed by τ.

Supply layer

The supply layer shares the node set V of the topological layer, but differs from it by being represented as a multigraph. A given pair of terminals may therefore be connected by multiple parallel supply edges, each corresponding to a scheduled service operated between them. This multigraph structure distinguishes the supply layer from the simple topological adjacency layer and enriches each terminal pair with service-level information, including the operator, departure window, service duration, and operating days. These attributes provide the basis for path enumeration.

Four index sets are used throughout for any edge on supply layer. First, $\Pi = \Pi^r \sqcup \Pi^a$ collects the rail operators or airlines active in the analysis window, $\mathcal{S}_\pi$ denotes the set of scheduled services run by provider $\pi \in \Pi$, $\mathscr{D}$ is the set of calendar days in the analysis window, and $\mathscr{T} = \{[0,1),[1,2),\dots,[23,24)\}$ partitions the day into 24 hourly departure windows. Writing $\sigma(s) = (\sigma_1(s),\dots,\sigma_{n_s}(s))$ for the ordered stop sequence of service s, a supply-edge is generated whenever a passenger can board on s at one station and alight at a later station of the same s without an intermediate transfer:

$$E^{\text{sup}} = \Big\{ (v_i, v_j, s) \,\Big|\, \pi \in \Pi,\ s \in \mathcal{S}_\pi,\ \exists k < \ell :\ \sigma_k(s) = v_i,\ \sigma_\ell(s) = v_j \Big\}. \tag{3.12}$$

The relaxation from $\ell = k+1$ to $k < \ell$ is what makes the supply layer a distinct layer from the structural layer rather than a relabelling of it. Under $\ell = k+1$ the two layers would share the same edge set and the supply layer would reduce to the structural layer with per-edge attributes. Under $k < \ell$ it captures the broader relation "two terminals reachable without alighting on the same service", which is the relation the path enumeration of Section 3.2.5 self-joins on. For instance, a Eurostar service running Paris-Gare du Nord → Brussel-Zuid → Amsterdam will be used to illustrate the difference (Figure 3.3). On the structural layer, only the two consecutive-stop edges Paris-Gare du Nord–Brussel-Zuid and Brussel-Zuid–Amsterdam will be recorded, while the supply layer adds a third edge Paris-Gare du Nord–Amsterdam for the passenger who boards in Paris and alights in Amsterdam without leaving the train. Forgetting the service index s and the ordering condition $k < \ell$ projects E^{sup} onto the structural edge set, so the topological layer is recovered as the simple-graph quotient of the supply-layer multigraph.

Each supply-edge $e = (v_i, v_j, s) \in E^{\text{sup}}$ carries the per-service attribute tuple

$$\alpha(e) = \big(\tau(e),\ \pi(e),\ t_{\text{dep}}(e),\ t_{\text{arr}}(e),\ M(e)\big), \tag{3.13}$$

where $\tau(e) \in \{r, a\}$ tags the mode of the underlying service, $\pi(e) \in \Pi$ identifies its operator, $t_{\text{dep}}(e)$ and $t_{\text{arr}}(e)$ are the scheduled departure and arrival times at v_i and v_j, and $M(e) \in$

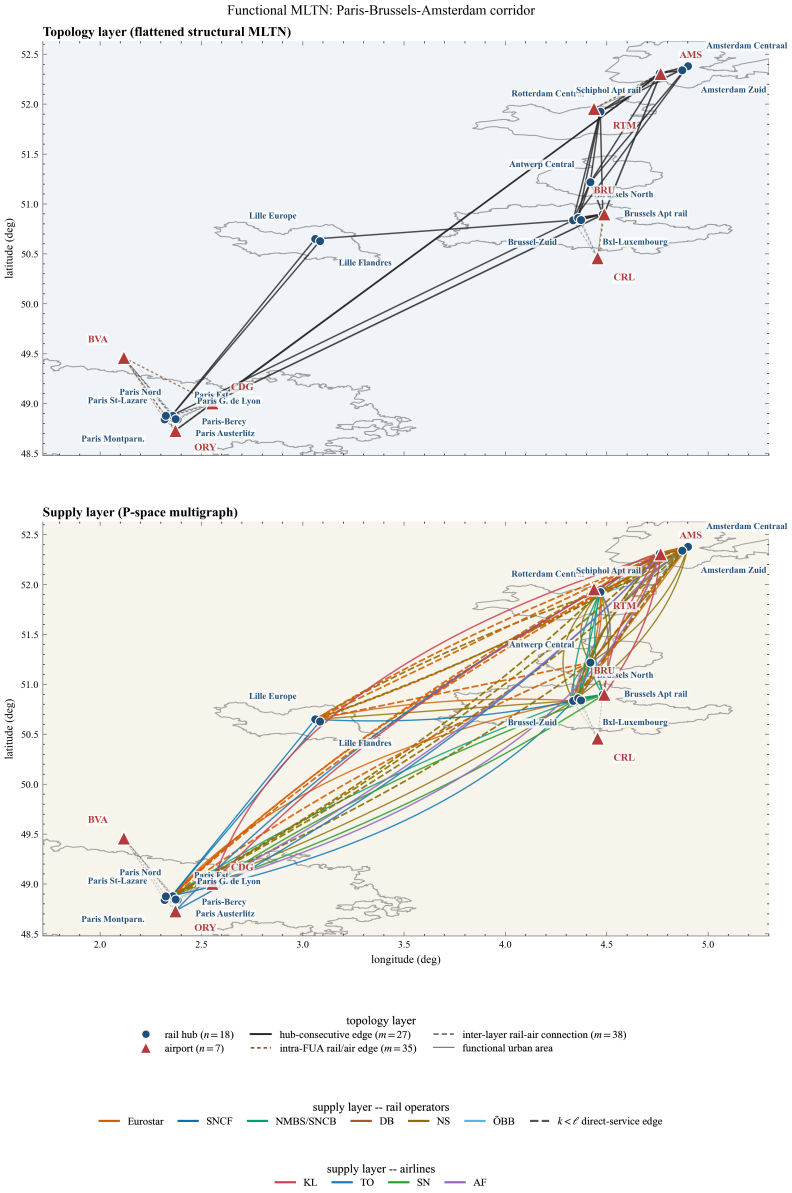

Figure 3.3: Supply-layered functional MLTN of the Paris-Brussels-Amsterdam corridor, spanning the Paris, Lille, Bruxelles/Brussel, Antwerpen, Rotterdam, and Amsterdam functional urban areas. Both panes share the same node set of eighteen intercity stations and seven airports. The topological layer (top pane) flattens the mode-layered structural MLTN of Section 3.1.3, retaining $m = 27$ consecutive-stop edges (23 rail and 4 air), $m = 35$ intra-FUA access edges, and $m = 38$ inter-layer rail-air footpath connections. The supply layer (bottom pane) replaces every topology edge by one parallel arc per scheduled service, giving the multigraph of equation 3.10. On the rail side this produces 73 service edges across 36 station pairs, coloured by operator (Eurostar, SNCF, NMBS/SNCB, DB, NS, OBB); dashed arcs mark the single skip-stop service whose endpoint pair carries no consecutive-stop edge in the topological layer. On the air side it produces 10 service edges across 4 airport pairs, coloured by carrier (KL, AF, TO, SN). Comparing the two panes makes the information added by the supply layer immediate, parallel operators on the same physical edge become visible, and service frequency emerges as a count of arcs per pair.

$\{0,1\}^{|\mathscr{D}|}$ is the operating-day bitmask of s. The in-vehicle time of the leg, $t^{\text{ivt}}(e) = t_{\text{arr}}(e) - t_{\text{dep}}(e)$, is a derived attribute that supplies one term of the path-level generalised travel cost of equation 3.21. The GTC function would then rely on the supply-edges enriched service information to form the equations.

For experiments that operate at terminal-pair granularity rather than supply-edge granularity (notably the supply-weighted centrality of Section 3.3.3), the parallel supply-edges between any pair (v_i, v_j) are summarised by four pair-level aggregates:

- $\Pi(v_i, v_j) = \{\pi(e) : e = (v_i, v_j, \cdot) \in E^{\text{sup}}\}$ — the set of providers offering at least one service between v_i and v_j, with per-provider multiplicity

$$n_\pi(v_i, v_j) = \big|\{e = (v_i, v_j, \cdot) \in E^{\text{sup}} : \pi(e) = \pi\}\big|. \tag{3.14}$$

  Eurostar and Thalys both contribute parallel supply-edges to the Paris–Brussels pair; British Airways, easyJet, and Air France all contribute parallel supply-edges to the London–Paris pair. Both cases are encoded by $|\Pi(v_i, v_j)| > 1$ and are preserved as parallel edges in the multigraph rather than collapsed into a single edge with an operator-set attribute;
- $n(v_i, v_j) = \sum_{\pi \in \Pi(v_i, v_j)} n_\pi(v_i, v_j)$ — the total number of parallel supply-edges between the pair, used as the supply-side edge weight in the experiments of Section 3.3.3;
- $M(v_i, v_j) = \bigvee_{e=(v_i, v_j, \cdot)} M(e)$ — the pair-level operating-day mask, defined as the bitwise OR of the per-service masks, with $M(v_i, v_j)_d = 1$ iff at least one service between the pair operates on calendar day d;
- $T(v_i, v_j) \in 2^{\mathscr{T}}$ — the pair-level departure-window set, $T(v_i, v_j) = \bigcup_e T(e)$, summarising the hourly slots in which at least one supply-edge between the pair departs.

The intra-FUA connection edges, namely E^r_{intra}, E^a_{intra}, and the interdependent coupling edges E_{inter}, are not part of the scheduled long-distance service supply. Instead, their travel-time summaries are queried on demand from the urban routing engine whenever a path-enumeration alternative requires a terminal change within an FUA. These times enter the GTC formulation in equation 3.21 through the connection-time term t^{connect}_τ.

Multiplex coupling

Because the functional MLTN G^{f} is a multiplex network in the strict sense of Šfiligoj et al. (2026, §2.2), the inter-layer block is the identity coupling between replicas of the same physical node:

$$a_{ij}^{\text{topo,sup}} = \delta_{ij}, \qquad i, j = 1, \ldots, |V|, \tag{3.15}$$

where δ_{ij} is the Kronecker delta, so each node on the topological layer is coupled to its own replica on the supply layer and to no other. The vertical coupling is therefore categorical, unweighted, and one-to-one, since every node $i \in V$ that carries scheduled service is present on the topological layer and mirrored by a single identical replica on the supply layer, the correspondence visible in the shared node set of Figure 3.3.

Supra-adjacency form.

The functional MLTN admits the block representation

$$\mathbf{A}^{\text{f}} = \begin{bmatrix} A^{\text{topo}} & I_{|V|} \\ I_{|V|} & A^{\text{sup}} \end{bmatrix}, \tag{3.16}$$

where $I_{|V|}$ is the $|V|\times|V|$ identity matrix encoding the multiplex coupling of equation 3.15. The diagonal block A^{topo} is the simple-graph adjacency of Section 3.1.3; the diagonal block A^{sup} is the multigraph adjacency of the supply layer, with $A^{\text{sup}}_{ij} = n(v_i, v_j)$ recording the number of parallel supply-edges between v_i and v_j. The per-supply-edge attributes of equation 3.13 are indexed separately and consulted when scenarios restrict to a particular operator (Section 3.4.2), calendar day, or hourly slice. The block structure matches Sfiligoj et al.'s canonical multiplex form (their figure 4b): square diagonal coupling between identical node replicas, no cross-coupling between distinct supply strands.

3.2. Modelling of travel time on the network

Services are performed by vehicles operating on physical infrastructure, but the operators and markets that provide European long-distance transport are fragmented on the supply layer. Rail remains largely shaped by national systems and is therefore best understood as a network of nationally framed infrastructure managers and incumbent operators (Cats, 2025; Pellegrini & Rodriguez, 2013). Air transport, by contrast, operates within a liberalised single market for carriers and aircraft, but is still constrained by national air-traffic-management blocks, slot-constrained hub airports, and hub-and-spoke carrier routing (Burghouwt, 2007; Pellegrini & Rodriguez, 2013). Accordingly both modes produce a similar passenger-facing consequence as long-distance journeys often require a change of operator, alliance, ticket, terminal, or mode, together with the associated transfer-time and time-window constraints.

Modelling pure vehicle-rotation patterns would therefore understate the effective connectivity of the European long-distance rail and air system, since the binding constraints on a passenger journey come from these operator-, slot-, and timetable-level interlocks rather than from the physical reach of a single train or aircraft. Modelling travel time from a passenger perspective and tracking it in relation to disruption scenarios therefore captures effective connectivity more faithfully. The passenger-oriented approach also makes explicit the transfer-time and time-window constraints that mediate cross-operator and cross-modal connectivity, and it could serve as the initial input for future flow-distribution models if passenger demand is subsequently integrated.

The study uses some assumptions from travel-behaviour and travel-demand research, but since the focus is on network service characteristics, the travel time is modelled from terminal to terminal, where a terminal is either a long-distance rail station or a commercial airport in the FUA-based universe of Section 3.1.1. This terminal-to-terminal segment is only one part of the door-to-door travel time. The travel time of interest therefore consists of two main components, namely in-vehicle time (IVT) spent aboard a train or aircraft, and out-of-vehicle time (OVT) spent waiting, walking, transferring between services (Ortúzar & Willumsen, 2011; Wardman et al., 2016). The empirical literature consistently finds that passengers perceive OVT as 1.5 to 2.5 times more onerous than IVT, with the exact ratio depending on the type of out-of-vehicle activity and the trip context (Wardman et al., 2016). At airports the OVT envelope additionally absorbs deplaning, baggage reclaim, immigration, security clearance, and boarding flows, which are materialised in the model as the mode-dependent arrival buffers and boarding floors of Section 3.2.3 rather than as separate IVT components.

The terminal-to-terminal scope is adopted because the study evaluates network performance, and the time passengers spend inside the transport network itself is therefore the primary object of measurement. By contrast, modelling home-to-terminal time would introduce large amounts of access and egress time that depend primarily on household car ownership and on local urban transit, neither of which is the object of this study. First-mile and last-mile connection trips are omitted for the same reason. Kristoffersson and Berglund (2022) observe that these connection segments account for 30% to 50% of total trip cost in air travel against under 10% for rail, because railway stations are typically more centrally located than airports. Omitting access and egress is therefore conservative for rail and lenient for air, so that any comparative advantage that air displays in the disruption experiments of Chapter 5 should be read as an upper bound on its real advantage once full door-to-door access and egress are reinstated. Intermediate processing inside the network boundary is kept in the model. At every station or airport reached along a multi-leg itinerary, mode-dependent processing buffers and platform-side waiting are computed explicitly so that air-involving transfers are not penalised by rail-rail conventions.

The IVT and OVT components are combined into a single generalised travel cost (GTC), which weights out-of-vehicle time more heavily than in-vehicle time to reflect the higher perceived disutility of waiting, walking, and the time spent on transfer connection. The following subsections describe each time component, the multiplier applied to it, and how the generalised travel cost is ultimately defined and computed from them.

3.2.1. In-vehicle travel time

The in-vehicle travel time of a leg (t_l^{ivt}) is computed directly from the published timetable of the operating service as the difference between the scheduled arrival time at the alighting facility and the scheduled departure time at the boarding facility. The definition applies to both modes, with two practical specifics. For a rail leg, the boarding and alighting facilities are any two stations along the same trip with the boarding stop preceding the alighting stop in the stop sequence, so a single long-distance service yields IVT values for many station pairs along its route. For an air leg, every scheduled service in the analysis window is a non-stop flight between an origin and a destination airport, and the difference between the scheduled arrival and departure times is the gate-to-gate block time, which already absorbs taxi-out, airborne, and taxi-in time by construction. In both cases t_l^{ivt} enters the generalised travel cost with a weight of 1.0 (the reference unit), against which the transfer-time multipliers are calibrated. For paths with multiple legs, the total in-vehicle travel time is the sum of the per-leg values across all legs, regardless of the mode mix.

3.2.2. Out-of-vehicle time and transfer modelling

The out-of-vehicle time (OVT) captures the time spent outside a vehicle during a journey that enters the generalised travel cost as priced minutes. The model decomposes OVT into two components, namely a waiting component at the boarding facility, and a movement component between two facilities within the same FUA when consecutive legs are not boarded at the same place. The two components are distinguished because they impose different types of disutility on passengers, and the empirical literature assigns them different multipliers relative to in-vehicle time (Wardman et al., 2016). The duration of

facility-internal processing flows such as deplaning, baggage reclaim, immigration, security clearance, and boarding is deliberately not priced as OVT, because it is per-passenger heterogeneous and is not observable from the schedule. The passenger is therefore treated as available to start the next-leg transfer the moment the arriving service reaches its alighting facility, and the mode-dependent arrival buffer and boarding floor of Section 3.2.3 act as a feasibility guard rather than as GTC components, ensuring that the wall-clock gap at any transfer is large enough to absorb the actual processing time and any service lateness. The generalised travel cost of a path is therefore an optimistic lower bound on the experienced passenger cost, and the buffer constants are calibrated so that the bound is achievable rather than aspirational.

Transfers arise from structurally different drivers in the two modes. On the rail side they reflect the nationally fragmented network already discussed in Section 3.2, where cross-border services are scarce and passengers must often change operator or service line (Calzada-Infante et al., 2020; European Court of Auditors, 2018; O'Sullivan & Patel, 2004). On the air side the hub-and-spoke organisation of European carriers means that peripheral airport pairs are reached only through one or two hub airports, even though the airspace itself is unified (Burghouwt, 2007). The model therefore restricts a journey to at most two transfers and constrains flight-to-flight changes to one, so that an itinerary carries no more than three legs and at most two of them by air. The one-air-transfer ceiling is also grounded in observed itinerary choice. By modelling air-travel itinerary choice with a level-of-service dimension that separates nonstop, direct, single-connect, and double-connect products, Coldren and Koppelman (2005) find that, relative to a nonstop service, the estimated penalty on a double connection is more than twice that on a single connection, so a second flight-to-flight change moves an itinerary into a far less attractive class. More recent itinerary-choice estimates on independent data confirm that this preference for fewer connections is stable over time and across booking channels (Lurkin et al., 2018). Restricting travel by air to a single transfer therefore reproduces the connecting itineraries passengers actually use, while the air-air-air chain excluded by this rule is the rare itinerary that is both scarce in the schedule and much less preferred and travelled by passengers.

The upper bound of two transfers on rail is imposed primarily for computational reasons. Allowing unrestricted transfers would otherwise generate highly implausible itineraries. The worst case binds harder on rail, where preserving connectivity for peripheral cross-border OD pairs can produce routes requiring as many as nine to eleven transfers, than on air, where even a pathological hub cascade less likely to involve more than two transfer. Such paths are unlikely to be chosen by passengers and therefore do not represent plausible travel behaviour. The ceiling is also supported by the stated-preference evidence of Curtale et al. (2023), who found in a rail context that the first transfer accounts for the largest share of perceived disutility, with the marginal penalty of a second transfer substantially smaller. The pattern of diminishing marginal disutility per additional transfer is a general behavioural feature that extends naturally to intermodal itineraries.

The two-transfer and one-air-transfer ceilings together fix the set of mode chains the enumerator can return. With each leg being rail or air and at most three legs, the number of ordered chains is $2^1 + 2^2 + 2^3 = 14$. The one-air-transfer rule removes air-air-air, the only chain with two consecutive flight-to-flight changes, leaving $14 - 1 = 13$ chains (Ta-

ble 3.1). They span the two direct services, the four single-transfer chains, and the seven two-transfer chains, among which rail-air-rail is the canonical intermodal itinerary that flies the long middle leg and uses rail for access and egress.

Table 3.1: The thirteen mode chains returned by the path enumeration, grouped by number of in-vehicle legs. The two-transfer ceiling caps a chain at three legs, and the one-air-transfer ceiling removes the single chain air-air-air, which would require two consecutive flight-to-flight changes.

Transfers	Legs	Mode chains
0	1	rail, air
1	2	rail-rail, rail-air, air-rail, air-air
2	3	rail-rail-rail, rail-rail-air, rail-air-rail, rail-air-air, air-rail-rail, air-rail-air, air-air-rail

Because transfers impose multiple forms of burden, including waiting as time passes, spatial effort associated with walking or local-transit navigation, and coordination effort arising from the risk of missing a connection, the model decomposes transfer time into two components, namely a waiting component at the boarding facility and a movement component between facilities. The two components carry different multipliers in the GTC of equation 3.21 and are described in the following two subsections. The mode mix at a transfer (rail-to-rail, rail-to-air, air-to-rail, air-to-air) does not alter this two-component decomposition but does change the magnitudes of the buffers and floors that bound the feasible transfer window, as set out in Section 3.2.3.

Same-terminal transfer (waiting)

A transfer is classified as same-terminal when the passenger can board the consecutive service directly at the facility where the arriving leg alights, without moving to a different station or airport. For same-terminal transfers the transfer can be treated as "extended waiting", since the passenger remains within the same station or airport and changes only platform or gate via dedicated walkways. The spatial cost relative to a transfer that requires boarding at a different facility is minimal, so the entire transfer-window time at the boarding facility is assigned to a single waiting component. This component is penalised with the waiting-time multiplier $\mu_{\text{wait}} = 1.76$, derived from the inter-urban meta-analysis of Wardman et al. (2016).

Inter-terminal transfer (movement between terminals)

A transfer is classified as inter-terminal when the passenger must move to another terminal because the consecutive service is not available at the arrival terminal. Four mode-pair cases are covered, namely rail-station to rail-station moves between two intercity stations in the same urban area, rail-station to airport moves for an onward flight, airport to rail-station moves for an onward train, and airport to airport moves between two airports in the same urban area. The movement leg between the alighting facility and the next boarding facility entails additional spatial effort and introduces coordination burden in the timing of consecutive trip legs. It is treated as the access leg of the following service and penalised with the access and egress multiplier $\nu_{\text{move}} = 1.87$, also derived from Wardman et al. (2016). The decoupling of the two multipliers is motivated by the observation that

walking or transferring with luggage between facilities imposes a higher disutility than waiting on a platform or in a gate area, consistent with the distinction made in WebTAG TAG Unit A1.3 (Department for Transport, 2024) between access and egress and waiting components.

As in previous studies, a complete machine-readable urban public-transport timetable is not available for every major European city (Brons et al., 2023). Although an OpenStreetMap file covering the European study area is available, enumerating all intra-FUA facility pairs across all FUAs that contain more than one terminal would be computationally prohibitive. To streamline data acquisition, the Google Directions Transit API is used to obtain intra-FUA transfer times for each of the four mode-pair cases, namely rail-rail, rail-airport, airport-rail, and airport-airport facility pairs.

3.2.3. Transfer window

Having established the disutility that transfers impose on passengers, the path-enumeration algorithms must define the feasible time window for a transfer connection to consecutive legs. Since rail and air transfers involve very different physical processes, the transfer window is specified as a set of mode-dependent constants rather than as a single global interval. For example, a rail transfer may only require a platform change within a station, whereas an air transfer may involve baggage reclaim, security clearance, and movement through a larger terminal environment. The rail-rail case is established first, since it inherits a substantial empirical literature, then the air-side constants are stated separately, and the two are combined into a unified feasibility constraint.

Rail-to-rail transfers

Rail-to-rail transfers include both same-terminal and inter-terminal transfers that connect two passenger legs within the rail mode. For these type of transfer, the lower bound on the platform-side waiting time is set at 5 minutes, which is the minimum time required for a platform change. The empirical findings of de Keizer et al. (2015) demonstrate, through a large-scale stated-preference study, that passengers perceive a 5-minute transfer as optimal. For transfer times shorter than 5 minutes, disutility rises sharply due to stress and the risk of missing the connection, while each minute beyond 5 minutes incurs an additional penalty equivalent to 1.67 in-vehicle minutes. This non-linear disutility structure, validated against real-world ridership data on Dutch intercity corridors, implies that 5 minutes is the transfer time at which the generalised travel cost associated with transferring is minimised.

The upper bound at any station is set at 60 minutes, consistent with the connection threshold adopted by Calzada-Infante et al. (2020) for European rail path enumeration and the computational constraint this study want to apply. The value reflects the operational reality that cross-border service frequency on international corridors is often limited to one or two departures per hour and a shorter bound would exclude feasible connections on these corridors. Major European intercity rail stations also function as multimodal hubs with commercial facilities, waiting lounges, and weather-protected concourses, which support planned waiting times of this length. A sensitivity analysis confirms that the 60-minute threshold sits at the transition from declining marginal connectivity gains to a plateau, and that varying the bound between 45 and 90 minutes does not qualitatively

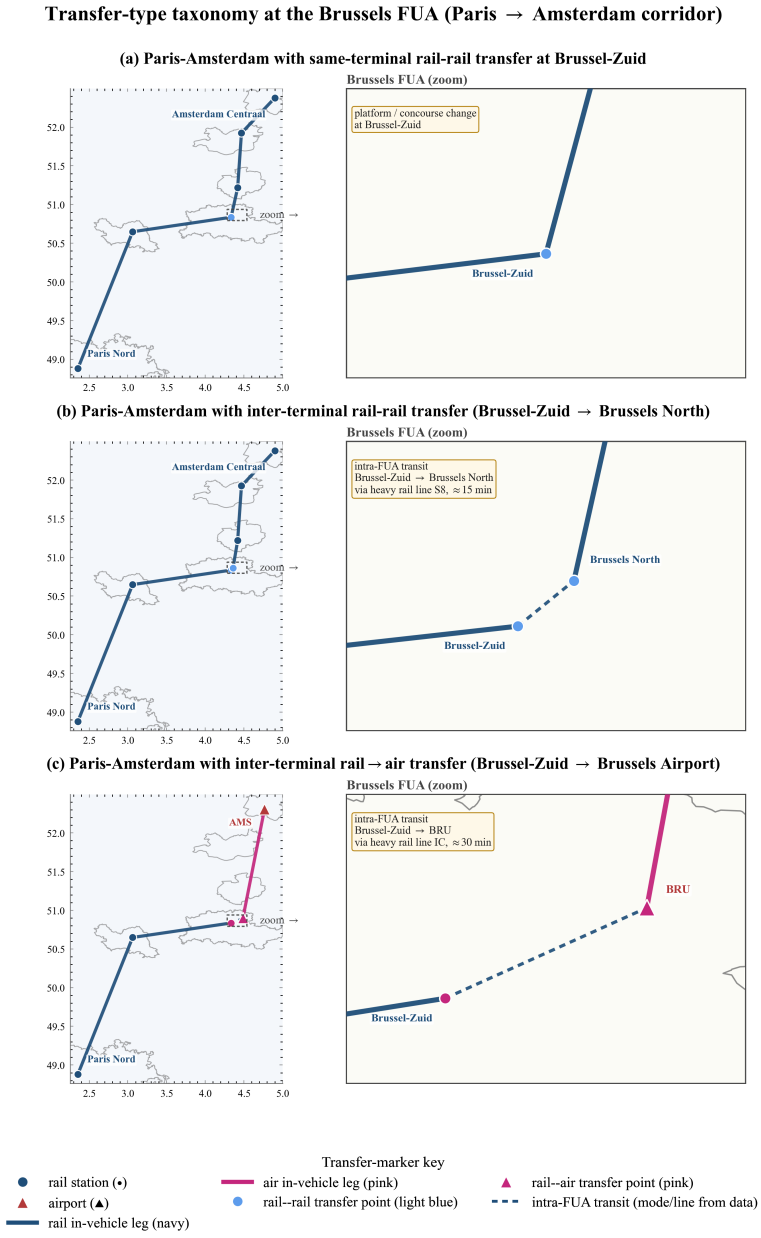

Figure 3.4: Three representative Paris → Amsterdam alternatives, each realising a different type of transfer at the Brussels FUA. Each row shows the full Paris → Amsterdam corridor on the left and a Brussels FUA close-up on the right. Panel (a) illustrates a same-terminal rail-rail transfer at Brussel-Zuid, where the passenger only changes platform within the same station. Panel (b) shows an inter-terminal rail-rail transfer between Brussel-Zuid and Brussels North, realised by an intra-FUA heavy-rail leg on line S8 of approximately 15 minutes. Panel (c) shows an inter-terminal rail-to-air transfer from Brussel-Zuid to Brussels Airport, realised by an intra-FUA heavy-rail leg on line IC of approximately 30 minutes, followed by a direct flight to Amsterdam Schiphol. Navy lines represent rail in-vehicle legs and pink lines represent air in-vehicle legs. Same-terminal transfer points are marked by a single circle, while inter-terminal transfer points are marked by a pair of facility-shape markers (circle for rail station, triangle for airport) connected by a dashed line whose colour and label reflect the actual transit mode obtained from the Google Directions Transit lookup of Section 3.2.2.

change the robustness results. On the arrival side, an additional buffer of $B_{\text{rail}} = 50$ minutes is applied between the scheduled arrival of a train and the moment the passenger is treated as available to board onward services; it absorbs alighting, exit from the platform, and orientation inside the station, and is calibrated against the IRTAD interchange-time figures collated by de Keizer et al. (2015).

Air-involving transfers

For transfers that arrive or depart by air, the arrival-side buffer and the boarding-side floor reflect the operational time costs of airport-side processing. Large hub airports, e.g. Charles de Gaulle, Heathrow, Frankfurt, Amsterdam Schiphol, and Madrid-Barajas require an arrival buffer of $B_{\text{air,large}} = 120$ minutes covering deplaning, baggage reclaim where applicable, immigration, and transit through the terminal, and a boarding floor of $W_{\text{min}}^{\text{air,large}} = 120$ minutes covering check-in, security, clearance, and boarding. Medium and small commercial airports operate with shorter values of $B_{\text{air,medsmall}} = W_{\text{min}}^{\text{air,medsmall}} = 90$ minutes, reflecting the lighter processing load of regional terminals. The 120-minute and 90-minute values, and the underlying binary partition into large and medium-small airports, follow Bruno et al. (2025), who adopt these figures to represent realistic average transfer behaviour at European airports in preference to the minimum connecting times of 45 to 90 minutes conventionally used in the air-network literature. The numerical threshold that decides which airports fall into the large class is set in this study at 10 million scheduled passengers per year, drawing on EUROSTAT's air-passenger-transport statistics (Eurostat, 2026) joined to the airport node set by IATA code, with the annual figure computed as the sum of the most recent 12 reported months within a 2024-01 to 2026-02 window to keep the sample comparable across airports with different reporting lags. Airports in UK, which EUROSTAT no longer covers after Brexit, are supplemented from the UK Civil Aviation Authority's "Size of UK Airports" table for the 2025 annual reference period (UK Civil Aviation Authority, 2026).

The boarding ceiling at any airport is set at the trip-elapsed cap $W_{\text{max}}^{\text{air}} = 1{,}440$ minutes. The air layer is far thinner than the rail layer in both coverage and time, spanning 197 airports against 1,674 rail stations and drawing its schedule from a single month of commercial timetables, whereas the rail feeds combine several timetable editions. A tighter ceiling would prune feasible air connections that exist only at long layovers and so would understate this already sparse connectivity, and the longest wait a traveller will accept at an airport has no clean empirical value. A deliberately loose ceiling is therefore adopted as a feasibility floor for thin coverage rather than as a behavioural assumption, and its effect on the enumerated alternatives is examined in the sensitivity analysis of Section 5.4.1.

Feasibility constraint

Let τ index a transfer event at facility k between the arriving leg l_{arr} and the departing leg l_{dep}, and write $\text{mode}(k) \in \{\text{rail}, \text{air,medsmall}, \text{air,large}\}$ for the mode-and-size class of facility k. Let t_τ^{gap} denote the elapsed time between the arrival of l_{arr} at its alighting facility and the scheduled departure of l_{dep}, and let t_τ^{connect} denote the intra-FUA connecting time between the two facilities (zero for same-terminal transfers). The arrival-side buffer $B_{\text{mode}(k)}$ and the boarding-side floor $W_{\text{min}}^{\text{mode}(k')}$ at the boarding facility k' are drawn from

Table 3.2. The transfer is feasible iff

$$B_{\mathrm{mode}(k)} + t_\tau^{\mathrm{connect}} + W_{\min}^{\mathrm{mode}(k')} \le t_\tau^{\mathrm{gap}} \le W_{\max}^{\mathrm{mode}(k')}. \tag{3.17}$$

The platform-side waiting time, i.e. the time the passenger spends at facility k' after the buffer and connection have been spent, is $t_\tau^{\mathrm{wait}} = t_\tau^{\mathrm{gap}} - B_{\mathrm{mode}(k)} - t_\tau^{\mathrm{connect}}$. The two time components enter the GTC with the multipliers introduced in Section 3.2.2: the connecting time at $\nu_{\mathrm{move}} = 1.87$ and the platform-side waiting time at $\mu_{\mathrm{wait}} = 1.76$.

Inter-terminal transfer

When the boarding facility k' differs from the alighting facility k (intra-FUA move), the transfer is feasible ff

$$B_{\mathrm{mode}(k)} + t_\tau^{\mathrm{connect}} + W_{\min}^{\mathrm{mode}(k')} \le t_\tau^{\mathrm{gap}}, \qquad t_\tau^{\mathrm{wait}} \le W_{\max}^{\mathrm{mode}(k')}, \tag{3.18}$$

with $t_\tau^{\mathrm{wait}} = t_\tau^{\mathrm{gap}} - B_{\mathrm{mode}(k)} - t_\tau^{\mathrm{connect}}$.

The two constants are read from different facilities. The arrival buffer $B_{\mathrm{mode}(k)}$ comes from the facility the passenger alights at and absorbs the incoming leg's processing and lateness, whereas the boarding floor and ceiling $W_{\min}^{\mathrm{mode}(k')}$ and $W_{\max}^{\mathrm{mode}(k')}$ come from the facility boarded next, so a rail-to-air transfer pairs the rail arrival buffer with the airport boarding floor and an air-to-rail transfer pairs the airport arrival buffer with the rail boarding floor, which is the source of the asymmetry between the two directions. Table 3.3 summarises, for each inter-terminal transfer type, the arrival buffer, connect time, boarding floor, and ceiling, and marks which components enter the generalised travel cost. The gap divides into the arrival buffer, the inter-terminal movement $t_\tau^{\mathrm{connect}}$, and the platform-side waiting t_τ^{wait} that is never shorter than the boarding floor, and of these only the movement and the waiting are carried into the generalised travel cost function in Section 3.2.4 while the buffer remains a feasibility requirement.

Same-terminal transfer

When $k = k'$ the passenger remains within the same station or airport, the connecting time vanishes, and the arrival buffer is absorbed into the planned waiting period rather than added on top. Feasibility reduces to

$$W_{\min}^{\mathrm{mode}(k')} \le t_\tau^{\mathrm{gap}} \le W_{\max}^{\mathrm{mode}(k')}, \tag{3.19}$$

which recovers rail-rail $[5, 60]$ transfer window from de Keizer et al. (2015) and Calzada-Infante et al. (2020) and the analogous $[120, 1440]$ window at a large hub airport.

Lastly, independently of the per-transfer feasibility constraint above, every multi-leg alternative is required to satisfy a 24-hour cap on total trip elapsed time:

$$t^{\mathrm{elapsed}}(p) = t_{l_{|p|}}^{\mathrm{arr}} - t_{l_1}^{\mathrm{dep}} \le 1440 \text{ min}, \tag{3.20}$$

which discards itineraries whose combined in-vehicle and out-of-vehicle time exceeds one day.

Table 3.2: Mode-dependent transfer-window constants used by the path enumeration.

Facility type	B_{mode}	$W_{\min}^{\text{mode}}$	$W_{\max}^{\text{mode}}$
Rail station	50 min	5 min	60 min
Airport, medium or small	90 min	90 min	1,440 min
Airport, large hub	120 min	120 min	1,440 min

Table 3.3: Cost components and feasibility bounds for each inter-terminal transfer type, in minutes. The arrival buffer is read from the alighting facility k and acts as a feasibility guard that is not priced, whereas the inter-terminal connect time and the platform-side waiting are priced at $\nu_{\text{move}} = 1.87$ and $\mu_{\text{wait}} = 1.76$. Large-hub airport values are listed first and medium or small airport values in brackets. The minimum feasible gap is reached when the platform-side wait equals the boarding floor $W_{\min}^{\text{mode}(k')}$.

Inter-terminal transfer	Arrival buffer $B_{\text{mode}(k)}$ (not priced)	Connect t_τ^{connect} (ν_{move})	Boarding floor $W_{\min}^{\text{mode}(k')}$ (μ_{wait})	Boarding ceiling $W_{\max}^{\text{mode}(k')}$	Minimum feasible gap
Station to station	50	t_τ^{connect}	5	60	$55 + t_\tau^{\text{connect}}$
Intermodal rail → air	50	t_τ^{connect}	120 [90]	1440	$170\,[140] + t_\tau^{\text{connect}}$
Intermodal air → rail	120 [90]	t_τ^{connect}	5	60	$125\,[95] + t_\tau^{\text{connect}}$
Airport to airport	120 [90]	t_τ^{connect}	120 [90]	1440	$240\,[180] + t_\tau^{\text{connect}}$

3.2.4. Generalised travel cost

The in-vehicle time of Section 3.2.1, the priced waiting and movement components of Section 3.2.2, and the feasibility-bounded transfer arithmetic of Section 3.2.3 are combined into a single passenger-perspective cost figure, the generalised travel cost (GTC). Each priced component enters with the multiplier introduced in its defining subsection and the GTC of a multi-leg path p is

$$\text{GTC}(p) = \sum_{l \in \mathscr{L}(p)} t_l^{\text{ivt}} + \mu_{\text{wait}}\Big(W^{\text{o}}(p) + \sum_{\tau \in \mathscr{T}(p)} t_\tau^{\text{wait}} + W^{\text{d}}(p)\Big) + \nu_{\text{move}} \sum_{\tau \in \mathscr{T}^{\text{move}}(p)} t_\tau^{\text{connect}}; \tag{3.21}$$

where $\mathscr{L}(p)$ is the set of in-vehicle legs along p, $W^{\text{o}}(p)$ and $W^{\text{d}}(p)$ are the origin and destination dwell times read from Table 3.4 according to the mode and hub class of the path's first and last terminals, $\mathscr{T}(p) = \mathscr{T}^{\text{same}}(p) \cup \mathscr{T}^{\text{move}}(p)$ is the transfer set with same-terminal and inter-terminal subsets, t_τ^{wait} is the platform-side waiting at transfer τ, and t_τ^{connect} is the intra-FUA movement at an inter-facility transfer. The cost separates into three groups.

The first term sums the in-vehicle time across every leg of p at the reference weight of 1.0. The units are commensurable across modes because both rail and air leg durations are measured in scheduled timetable minutes per Section 3.2.1, so a multi-leg sum across a rail-and-air itinerary is well-defined without further conversion.

The second term applies the waiting-time coefficient, $\mu_{\text{wait}} = 1.76$, to all platform-side waiting, whether at a same-terminal or inter-terminal transfer, and to the origin and destination dwell. At a same-terminal transfer the passenger remains within one station or airport, the arrival buffer collapses to zero per the special-case rule of Section 3.2.3, and

Table 3.4: Origin and destination terminal dwell priced in the generalised travel cost, by mode and hub class. These cover the fixed processing a traveller incurs at the first and last terminal and are distinct from the mid-trip transfer buffers of Table 3.2.

Facility type	W^{o} (origin)	W^{d} (destination)
Rail station	25 min	5 min
Airport, medium or small	90 min	30 min
Airport, large hub	120 min	15 min

the full wall-clock gap is platform-side waiting that enters the GTC as priced minutes.

Pricing a fixed origin and destination dwell, rather than only in-vehicle and transfer time, is what keeps the cross-modal comparison fair. Air imposes a substantial fixed time cost at the first and last terminals, namely check-in, security, and boarding before departure and disembarking and egress after arrival, that rail largely does not. Were this cost omitted, a short flight would dominate a rail itinerary of comparable in-vehicle time on generalised cost alone, and the cross-modal comparisons and the substitution and rescue results of Chapter 5 would be biased towards air. Pricing the dwell, calibrated to the average European airport figures of Bruno et al. (2025), offsets the in-vehicle advantage of a flight by the terminal overhead a passenger actually incurs. It also partially counterbalances the omission of door-to-terminal access in Section 3.2, which on its own is lenient towards air, so that the two simplifications act in opposite directions rather than compounding into a single-sided bias towards either mode.

The third term prices inter-terminal transfer time at $\nu_{\text{move}} = 1.87$. At an inter-terminal transfer the priced component is the intra-FUA connecting time $t_{\tau}^{\text{connect}}$, drawn from the Google Directions Transit lookup of Section 3.2.2 for the relevant mode-pair (rail-rail, rail-air, air-rail, or air-air). The arrival buffer is absorbed as a feasibility guard rather than charged to the passenger.

The two multipliers $(\mu_{\text{wait}}, \nu_{\text{move}})$ are taken from the European inter-urban meta-analysis of Wardman et al. (2016), the standard empirical anchor for long-distance value-of-time research in this study area. A path with zero transfers carries no transfer waiting and no inter-terminal movement, so its generalised travel cost is the in-vehicle time plus the origin and destination dwell W^{o} and W^{d}. For a direct flight this terminal dwell is substantial, which is why a non-stop air service does not collapse to its block time alone.

The generalised travel cost does not charge the connection buffer, which is the fixed block of time set aside at each change for getting off the train or aircraft, reclaiming baggage, clearing security again, walking out, and covering the risk that the arriving service runs late. That block is real time, and part of it, the walking and the delay margin, is exactly the kind of time the value-of-time literature treats as onerous (Wardman et al., 2016), while connection time as a whole is known to make up a large share of an air trip (Kristoffersson & Berglund, 2022). The timetable does not reveal how the buffer divides between these uses, so the model does not try to weight it. It counts the buffer only toward the journey's elapsed time and toward whether the connection is possible at all, in the same way that Bruno et al. (2025) and Calzada-Infante et al. (2020) treat connection times as feasibility requirements rather than as a priced cost. The effect is that the generalised travel cost is a floor on what a traveller actually experiences, and that floor is loosest for itineraries with

the most changes, which are the intermodal ones. Pricing the buffer would raise the cost of air and intermodal itineraries in particular, so leaving it out makes the robustness and substitution findings cautious rather than flattering. Resolving the buffer into its walking, processing, and delay components and weighting each is left to future work.

3.2.5. Timetable-based path-finding algorithm

The travel-time modelling described in Section 3.2 quantifies passenger travel times on the network. The next step is to apply the path-finding algorithm to generate feasible paths representing individual journeys, subject to the constraints imposed by the transfer model. For robustness analysis, path redundancy should be assessed not only in terms of the number of available paths, but also in terms of the set of viable alternatives that remain under disruption. This makes it possible to capture the dynamic effect of disruption. When one path fails, how many alternatives remain available, and at what additional cost. As argued by Bešinović (2020), such effects are not captured by purely topological analysis.

In general, the algorithm enumerates feasible OD paths by iteratively composing direct service legs at transfer stations. Starting from the set of all direct (zero-transfer) connections derived from the rail GTFS timetable and the commercial air schedule, each subsequent stage extends the path set by joining existing paths with direct legs at stations or airports where a transfer is feasible, adding one additional transfer per stage. At each stage, candidate paths are filtered by transfer-time feasibility, date-range overlap, and inter-FUA origin-destination constraints, then reduced to a generalised-travel-cost band that retains every alternative within a fixed margin of the cheapest option for its origin, destination and departure window rather than a fixed number of alternatives per pair.. The enumeration proceeds through three stages, corresponding to paths with zero, one, and two transfers respectively (Figure 3.7).

The staged composition, direct services first, then one-transfer and two-transfer paths, mirrors the round-based structure of RAPTOR (Delling et al., 2015), in which round k resolves the journeys that use at most k trips, while the leg-by-leg composition without a shortest-path priority queue is the connection-scan idea of CSA (Dibbelt et al., 2018). Both algorithms return a single earliest-arrival or Pareto-optimal journey, whereas this study retains a whole band of alternatives around each local optimum. The unit of that retention is the bucket. A bucket groups every enumerated alternative that shares one origin, one destination, and one of the 24 hourly departure windows, where the origin and the destination are each a station or an airport. The bucket minimum is the lowest generalised travel cost among the alternatives in a bucket, and the procedure keeps every alternative whose generalised travel cost lies within a band above that minimum, 60 minutes for rail-only iterations and 180 minutes for air-involving iterations, so that connectivity and feasibility are not underestimated when the optimal path is later removed under disruption. The single-optimum efficiency of CSA and RAPTOR is traded for a complete feasible alternative set, which is the object a robustness analysis requires.

Prior to the model running, a strict same-day constraint requires the rail legs of a path to share at least one calendar day on which they all run. This matters because the rail timetable is assembled from dozens of national GTFS feeds whose validity periods differ and only partly overlap across the 2024 to 2026 analysis window, so joining legs from feeds that never coincide on a single day would produce a path that exists on no real day

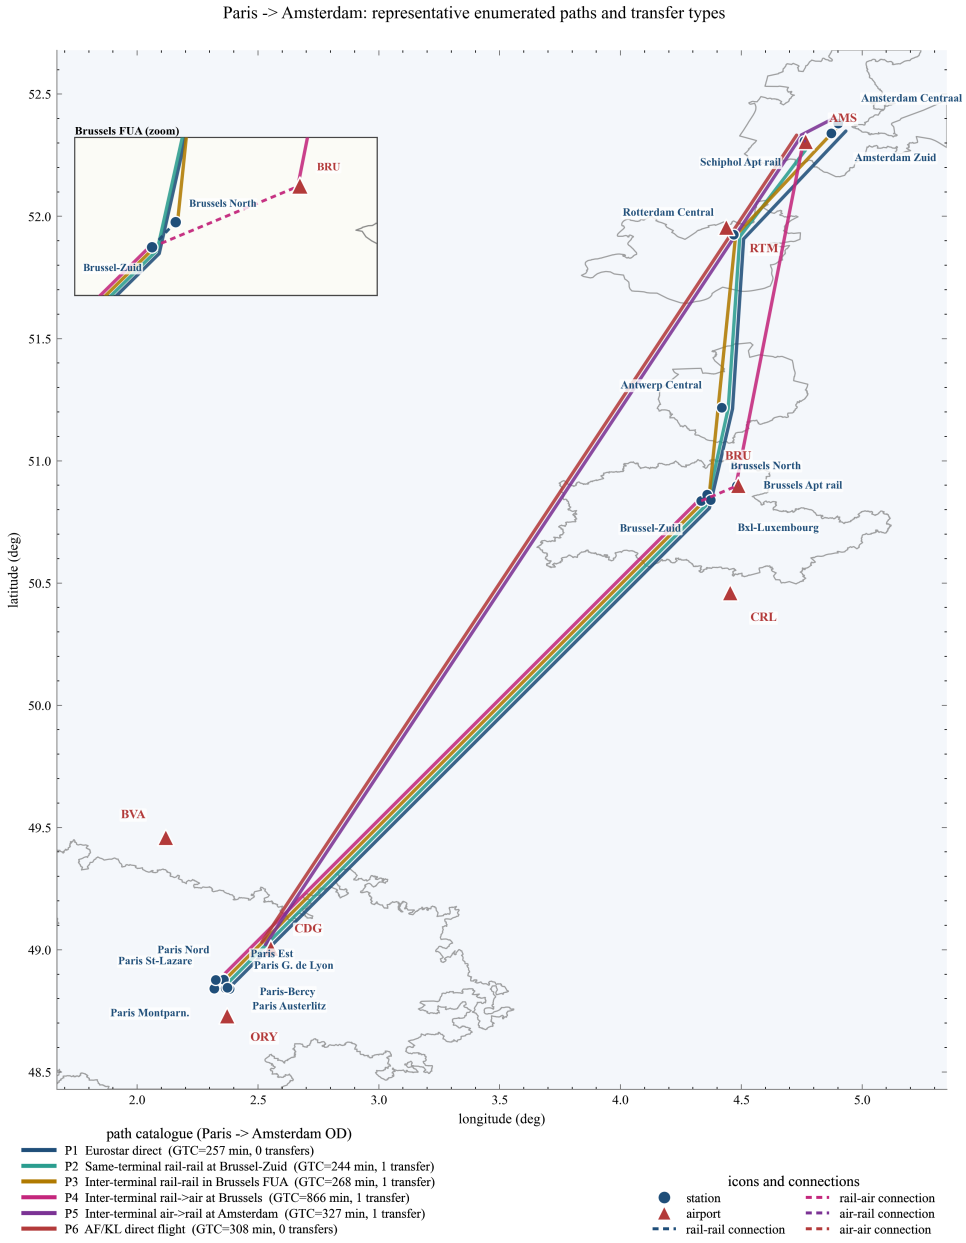

Figure 3.5: Six representative alternatives between the Paris and Amsterdam functional urban areas, chosen so that together they cover the full transfer taxonomy of sections 3.2.2 and 3.2.3, spanning a direct rail service and a direct flight, same-terminal and inter-terminal rail-rail changes, and the rail-to-air and air-to-rail transfers that cross between the two modes. Each inter-terminal transfer is drawn as a dashed connection line in the colour of its mode pair, and the Brussels inset enlarges three of them inside one FUA. The all-rail alternatives stay within a narrow range, from 243.9 to 268.0 minutes, with the small spread reflecting both which station they reach in the Amsterdam area and whether they run direct or change trains once. Adding an airport to the chain raises the generalised travel cost sharply, up to 866.5 minutes (more than fourteen hours) for the rail-to-air case, because each airport contact adds substantial waiting and connecting time. Intermodal travel, rather than the geographic distance, is what drives the cost a passenger faces.

(Figure 3.6). The check is made day by day rather than by comparing the start and end of each service's date range. In implementation, each service stores its operating days as a bitmask over the analysis window, and two legs are kept only when their masks share at least one common day, so a path survives only when its legs genuinely run together on some date. It should be noted that no single calendar date observes the full 30-country network at once. A comparison of the network across calendar time would therefore contrast different subsets of operators rather than one network at two times, and it would track data coverage rather than any property of the network. The network carried into the robustness analysis is for this reason the aggregate over the window, and its variation across calendar time is not estimated.

The air schedule is treated differently from the rail timetable in its calendar. The rail timetable is assembled from dozens of national feeds over a two-year GTFS feed coverage period, whereas the commercial air schedule covers a single representative month, March 2025. Matching a one-month air sample day for day against multi-year rail calendars is not meaningful, so each air leg is assumed to operate on every operating day of the rail legs it connects. The operating calendar of an intermodal itinerary is therefore the operating-date intersection of its rail legs alone, and an air leg adds no calendar restriction of its own. A relation served entirely by air is enumerated separately on the air schedule's own one-month calendar.

A further model simplification concerns the intra-urban connecting times. These are obtained from the Google Directions Transit API, queried at three representative departure times (07:00, 13:00, 22:00) to capture the time-dependent nature of urban public transport. The three departure windows are mapped to the 24 hourly departure windows of the timetable via a nearest-window rule with hours 0–9 use the morning window, hours 10–17 use the midday window, and hours 18–23 use the evening window. This trade-off reduces the number of API calls while still distinguishing between peak, off-peak, and evening urban transit conditions.

In the initial stage, the algorithm searches for direct inter-city paths that involve no transfers. For each rail trip in the GTFS timetable, every pair of stations where the origin departs before the destination arrives is recorded as a direct path, and every scheduled flight contributes the single non-stop airport pair it serves. Paths are assigned to one of 24 hourly departure windows ($\text{dep\_win} \in \{0, \ldots, 23\}$). Only inter-FUA paths are retained. Intra-FUA trips are filtered out to avoid confounding local and long-distance connectivity. Paths whose GTC exceeds the minimum for their OD pair by more than 60 minutes are discarded.

To construct paths involving one transfer, the model proceeds to the second stage. The unit of composition is a transfer facility, which may be a rail station or an airport. For each facility, every transfer-free leg arriving there is combined with every transfer-free leg departing from the same facility or from another facility in the same FUA reachable by an intra-urban connection, whose connecting time is drawn from the lookup of Section 3.2.2. Because those intra-FUA connections link the rail stations and airports that share an FUA, this single self-join produces all four transfer types at once, station to station, rail to air, air to rail, and airport to airport, with no mode-specific branching in the procedure. A candidate one-transfer path is kept only if its two legs use different routes, if the gap between arrival and onward departure satisfies the mode-dependent transfer

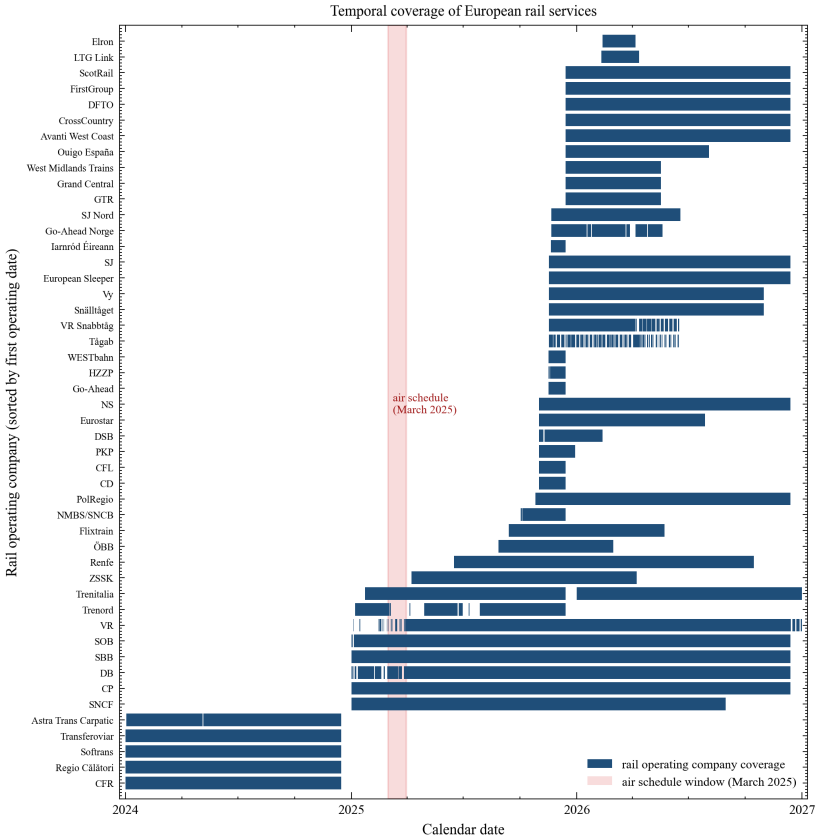

Figure 3.6: Temporal coverage of the rail services of the European rail operating companies. Each horizontal bar is one rail operating company, drawn over the calendar dates on which its services run and sorted by first operating date. Of the 51 rail operating companies in the operator universe, the 48 that carry an operating calendar in the timetable are shown here. The remaining three appear in the operator and route metadata but carry no scheduled trips, and so have no calendar to display. The rail service provided by different companies often start and end at different times and only partly overlap, which is why a path is kept only when its rail legs share at least one common operating day. The shaded band marks the single month of commercial air schedule (March 2025). Because that month is far narrower than the multi-year rail coverage, each air leg is assumed available on every rail-operating day rather than matched to the rail calendar day by day.

window of Section 3.2.3, whose constants in Table 3.2 pair the arrival buffer from the alighting facility with the boarding floor and ceiling from the facility boarded next, if the origin and destination lie in different FUAs under the membership map of equation 3.2, and if the rail legs share at least one common operating day. This last condition is the strict same-day check introduced above and illustrated in Figure 3.6, verified leg by leg on the operating-day masks of the rail legs, with an air leg assumed available on every rail-operating day so that the itinerary calendar is set by its rail legs alone (the calendar predicate of Algorithm 1).

The final stage concerns paths involving two transfers. The same composition logic combines each one-transfer path with one further transfer-free leg at a feasible facility, again spanning both modes. Combinatorial growth is held in check by the cost band itself, so a partial path is carried into this stage only while its generalised travel cost stays within the band, and the composed two-transfer paths are reduced to the same band afterwards. A further geographic constraint requires the FUAs of the origin, both transfer points, and the destination to be mutually distinct, which prevents circular or backtracking itineraries. The same-day check is applied between the new leg and the first leg of the existing path, since the internal consistency of the one-transfer path was established in the preceding stage. Figure 3.7 traces one retained two-transfer itinerary through these stages.

Algorithm 1 states the full procedure in the notation of this chapter, from the direct-leg base Π_0 through the one-transfer and two-transfer compositions Π_1 and Π_2 to the per-bucket band reduction, with the transfer-feasibility test of equation 3.17 factored out as the predicate Feasible and applied at every composition.

Algorithm 1: Timetable-based path enumeration as a staged self-join, in the notation of Chapter 3. Direct service legs are composed stage by stage with one further direct leg at a feasible transfer, and each stage is reduced to the generalised-travel-cost band above its per-bucket minimum.

Require: rail timetable $\mathcal{T}_r$, air timetable $\mathcal{T}_a$, intra-FUA connection lookup F, transfer constants $B, W_{\min}, W_{\max}$, functional-urban-area map $\mathrm{fua}(\cdot)$, band Δ_{band}, elapsed cap $C_{\max}$
Ensure: feasible alternative pool $\Pi_0 \cup \Pi_1 \cup \Pi_2$

Stage 1: direct legs (zero transfers). $\Pi_0 \leftarrow \varnothing$
for all trips $t \in \mathcal{T}_r \cup \mathcal{T}_a$ **do**
  for all stop pairs (s, s') of t with $\mathrm{dep}(s) < \mathrm{arr}(s')$ and $\mathrm{fua}(s) \neq \mathrm{fua}(s')$ **do**
    $\ell \leftarrow (s, s')$; $\mathrm{dep\_win}(\ell) \leftarrow \lfloor \mathrm{dep}(s)/60 \rfloor \bmod 24$; $\Pi_0 \leftarrow \Pi_0 \cup \{\ell\}$
  end for
end for
$\Pi_0 \leftarrow$ BandReduce(Π_0)

Stage 2: one transfer. $\Pi_1 \leftarrow \varnothing$
for all $p \in \Pi_0$ alighting at facility k **do**
  for all $\ell \in \Pi_0$ boarding at facility k' with $\mathrm{fua}(k') = \mathrm{fua}(k)$ **do**
    if Feasible(p, ℓ) **then**
      $\Pi_1 \leftarrow \Pi_1 \cup \{ p \oplus \ell \}$
    end if
  end for
end for
$\Pi_1 \leftarrow$ BandReduce(Π_1)

Stage 3: two transfers. $\Pi_2 \leftarrow \emptyset$ ▷ compose the band-reduced Π_1 directly
for all $p \in \Pi_1$ alighting at facility k **do**
  for all $\ell \in \Pi_0$ boarding at facility k' with $\mathrm{fua}(k') = \mathrm{fua}(k)$ **do**
    if Feasible(p, ℓ) **and** $\mathrm{fua}(\mathrm{src}\, p), \mathrm{fua}(k), \mathrm{fua}(k'), \mathrm{fua}(\mathrm{trg}\, \ell)$ all distinct **then**
      $\Pi_2 \leftarrow \Pi_2 \cup \{ p \oplus \ell \}$
    end if
  end for
end for
$\Pi_2 \leftarrow$ BandReduce(Π_2)
return $\Pi_0 \cup \Pi_1 \cup \Pi_2$

function Feasible(p, ℓ) ▷ transfer alighting at k, boarding at k'
  $g \leftarrow \mathrm{dep}(\ell) - \mathrm{arr}(p)$ ▷ gap in absolute minutes
  $t_{\mathrm{connect}} \leftarrow F(k, k', \mathrm{voronoi\_win}(\mathrm{arr}\, p))$ ▷ 0 when $k' = k$
  routes $\leftarrow \big(\mathrm{route}(\ell) \neq \mathrm{route}(\text{last leg of } p)\big)$
  window $\leftarrow \big(\mathbf{1}[k' \neq k]\, B_{\mathrm{mode}(k)} + t_{\mathrm{connect}} + W_{\min}^{\mathrm{mode}(k')} \leq g \leq W_{\max}^{\mathrm{mode}(k')}\big)$
  calendar $\leftarrow \big(\mathrm{cal}(p) \cap \mathrm{cal}(\ell) \neq \emptyset\big)$ ▷ op-mask bitwise AND $\neq 0$; an air leg's cal is all rail-operating days
  ends $\leftarrow \big(\mathrm{fua}(\mathrm{src}\, p) \neq \mathrm{fua}(\mathrm{trg}\, \ell)\big)$
  elapsed $\leftarrow \big(\mathrm{arr}(\ell) - \mathrm{dep}(p) \leq C_{\max}\big)$
  return *routes* $\wedge$ *window* $\wedge$ *calendar* $\wedge$ *ends* $\wedge$ *elapsed*
end function

function BandReduce(P) ▷ bucket $b = (\mathrm{src}, \mathrm{trg}, \mathrm{dep\_win})$
  return $\{ p \in P : \ \mathrm{GTC}(p) \leq \min_{p' \in b(p)} \mathrm{GTC}(p') + \Delta_{\mathrm{band}} \}$ ▷ $\Delta_{\mathrm{band}} = 60$ rail-only, 180 air-involving
end function

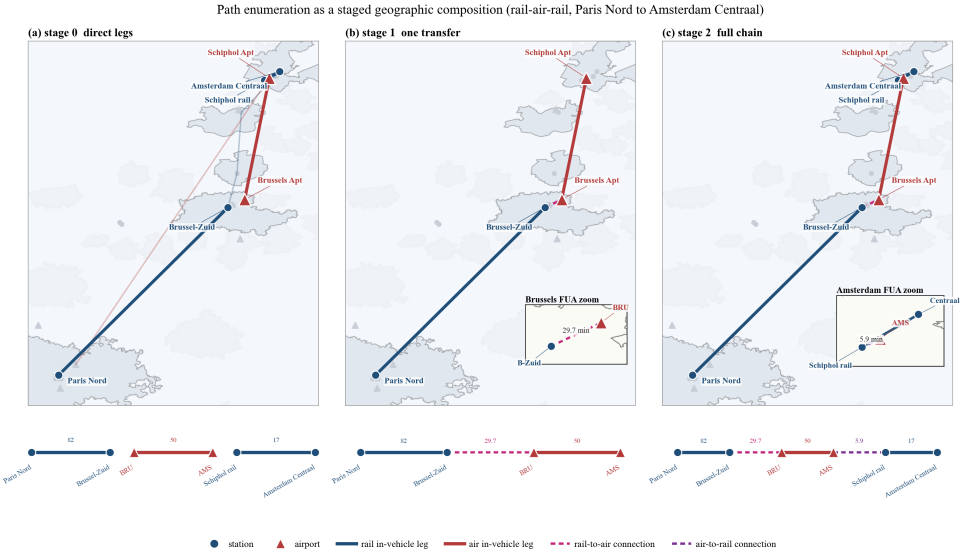

Figure 3.7: The path-enumeration process, illustrated by a rail-air-rail itinerary between Paris Nord and Amsterdam Centraal. The three panels share one map extent and add one stage at a time, and a node sequence diagram beneath each panel sets out the stations and airports in the order they are reached, with the time on each leg. Panel (a) is stage 1, the direct service legs retrieved from the rail and air timetables, with the intra-FUA connection lookup supplying the connecting times and the operating-day calendars the same-day feasibility check, and the faint legs marking the wider pool of direct services from which the example is composed. Panel (b) is stage 2, one transfer, joining the rail leg into Brussel-Zuid to the flight at the rail-to-air transfer enlarged in the Brussels inset, kept only when the gap covers the 50-minute arrival buffer, the 29.7-minute connecting time, and the 120-minute airport boarding floor, and when the rail and air legs share an operating day. Panel (c) is stage 3, the full chain, adding the air-to-rail transfer at Schiphol and the rail leg into Amsterdam Centraal. Its generalised travel cost follows equation 3.21, $\mathrm{GTC} = \sum_l t_l^{\mathrm{ivt}} + \mu_{\mathrm{wait}}\left(W^{\mathrm{o}} + \sum_\tau t_\tau^{\mathrm{wait}} + W^{\mathrm{d}}\right) + \nu_{\mathrm{move}} \sum_\tau t_\tau^{\mathrm{connect}} = (82 + 50 + 17) + 1.76\,(25 + 120 + 5 + 5) + 1.87\,(29.7 + 5.9) = 488.4$ minutes, with the endpoint and boarding dwells drawn from Table 3.4 and the arrival buffers entering only as a feasibility guard. The best alternative in this bucket is the direct rail itinerary at 256.8 minutes, and the air-involving band keeps every alternative within 180 minutes of that minimum, so this rail-air-rail itinerary, at 488.4 minutes, lies beyond the 436.8-minute band and is not eventually retained.

3.3. Robustness assessment framework

The robustness assessment framework employs metrics at two levels, including service-level metrics that incorporate generalised travel cost and robustness indices that summarise the degradation behaviour over the full range of disruption severity. Throughout, the unit of analysis is an origin destination pair paired with the hour of departure, written b and termed a bucket. The hour is kept because a connection that exists in the morning need not exist at night, so collapsing a whole day onto one figure per pair would let a single well-timed itinerary represent departures that have no comparable option. Each bucket is scored on its own minimum-cost itinerary, and the reported robustness measures then average over all buckets, so the results describe average connectivity across pairs and departure hours rather than a breakdown by time of day. The plain per-pair measures introduced first in Section 3.3.1 are a simpler starting point that the bucket-level

measures used for the results refine.

Every disruption scenario in this chapter, whatever its mechanism, is expressed as a single object, namely the set of enumerated alternatives that survive it, and every robustness indicator is a functional of that set. Let $\mathscr{P}$ be the enumerated alternative set of Section 3.2.5, partitioned into the departure buckets $\mathscr{B}$ of equation 3.36. Each alternative $p \in \mathscr{P}$ carries its generalised travel cost $g(p)$, its operator set $\Pi_p \subseteq \Pi$, and the structural edges $\Sigma(p) \subseteq \mathscr{U}^{\text{edge}}$ that it traverses. A disruption s at intensity x removes a set $R_s(x)$ from the targeted universe and leaves the surviving-alternative set

$$\mathscr{P}_s(x) = \{p \in \mathscr{P} \ : \ p \text{ uses no element of } R_s(x)\}, \tag{3.22}$$

where the meaning of "uses" is fixed per scenario in Sections 3.4.1–3.4.3. The surviving cost of a bucket is the minimum over its surviving alternatives,

$$g^{(x)}_{\min}(b) = \min\{g(p) : p \in \mathscr{P}_s(x),\ p \in b\}, \qquad g^{(x)}_{\min}(b) = +\infty \text{ if the bucket has no survivor.} \tag{3.23}$$

The robustness is therefore read on the surviving feasible itineraries, not on the connectivity of the structural graph, and a bucket counts as disconnected exactly when no enumerated alternative survives.

3.3.1. Service-level metrics

Service-level metrics incorporate the generalised travel cost computed through the path enumeration, enabling the assessment of disruption impacts on passenger connectivity rather than on abstract graph structure alone.

The *GTC-based efficiency* measures the mean reciprocal generalised travel cost across all buckets:

$$E_G(x) = \frac{1}{|\mathscr{B}|} \sum_{b \in \mathscr{B}} \frac{1}{\mathrm{GTC}^{(x)}_{\min}(b)} \tag{3.24}$$

where $\mathrm{GTC}^{(x)}_{\min}(b)$ is the minimum generalised travel cost for each bucket under the current network state, and $|\mathscr{B}|$ is the number of surviving buckets.

The *retained efficiency ratio* $R(x)$ expresses the post-disruption efficiency as a fraction of the baseline:

$$R(x) = \frac{E_G(x)}{E_G(0)} \tag{3.25}$$

where x is the disruption fraction (share of removed elements) and $E_G(0)$ is the baseline efficiency.

The *mean additional travel time* $\Delta\bar{D}(x)$ captures the average rerouting penalty:

$$\Delta\bar{D}(x) = \frac{1}{|\mathscr{B}|} \sum_{b \in \mathscr{B}} \min\big(\mathrm{GTC}^{(x)}_{\min}(b) - \mathrm{GTC}^{(0)}_{\min}(b),\ T(b)\big). \tag{3.26}$$

where the superscript (x) denotes the disrupted state and (0) the baseline. Buckets of station / airport pairs that become disconnected are assigned the mode-dependent reroute tolerance $T(b)$ of the operational definition in Section 3.3.3 (equation 3.39)

The *fraction disconnected* $f_{\text{disc}}(x)$ is the share of buckets for which no feasible path survives under the disruption scenario.

$$f_{\text{disc}}(x) = \frac{1}{|\mathscr{B}|}\left|\{b \in \mathscr{B} : \text{GTC}_{\min}^{(x)}(b) = +\infty\}\right| \tag{3.27}$$

3.3.2. Robustness indices

Robustness indices summarise the degradation curve over the full disruption range, following the taxonomy of Cats et al. (2017).

The *cumulative robustness-loss integral* F is the area above the retained-ratio curve:

$$F = \int_0^1 \left[1 - R(x)\right] dx \tag{3.28}$$

Values of F close to 0 indicate a robust network (the retained ratio stays near 1.0 across all disruption levels), while values close to 1 indicate rapid collapse. This metric corresponds to the cumulative fragility index used by Massobrio and Cats (2024) for metro networks.

The *accumulated delay* Q is the area under the mean-additional-travel-time curve:

$$Q = \int_0^1 \Delta\bar{D}(x)\, dx \tag{3.29}$$

This index captures the total passenger time penalty across all disruption levels.

3.3.3. Disruption operator and targeted-attack protocol

Robustness experiments are evaluated as a controlled sweep of disruptions of increasing intensity. To support both the topological scenarios on the structural MLTN of Section 3.1.3 and the service-level scenarios on the supply layer of the functional MLTN, the disruption is formalised as a parametric operator that returns a degraded network for any (intensity, targeting-rule) pair. The Type I to Type III scenarios of Section 3.4 are each obtained by selecting a particular targeting rule and applying the same operator and the same metric battery to the result.

Disruption operator

Let $G \in \{G^{\text{s}}, G^{\text{f}}\}$ denote either the structural or the functional MLTN. A disruption scenario is a map

$$\Phi_s : G \times [0,1] \longrightarrow G, \qquad (G, x) \longmapsto G_s(x), \tag{3.30}$$

where $x \in [0,1]$ is the disruption intensity, defined as the fraction of the targeted element universe that is rendered inoperable, and $s \in \mathscr{S}$ is the targeting rule drawn from the catalogue

$$\mathscr{S} = \{\text{rand}, \deg^{\text{topo}}, b^{\text{topo}}, \hat{b}, \text{operator}\}. \tag{3.31}$$

The element universe is fixed before the sweep: for granularity node, $\mathscr{U} = V$ (the rail-station or airport node set); for granularity edge, the universe is the long-distance subset of the structural edge set,

$$\mathscr{U}^{\text{edge}} = E^{r}_{\text{intra,ld}} \cup E^{a}_{\text{intra,ld}} \cup E_{\text{inter}}, \tag{3.32}$$

i.e. the inter-FUA service segments of equation 3.6 together with the cross-mode FUA-access edges of equation 3.7. Intra-FUA same-mode edges $E^{r}_{\text{intra,loc}} \cup E^{a}_{\text{intra,loc}}$ are excluded from $\mathscr{U}^{\text{edge}}$, since the chapter's scope is the European intercity system rather than metropolitan transit. Writing $K(x) = \lceil x|\mathscr{U}| \rceil$ and $R_s(x) \subseteq \mathscr{U}$ for the removed set under rule s at intensity x, the degraded network is

$$G_s(x) = G \setminus R_s(x), \qquad |R_s(x)| = K(x), \tag{3.33}$$

where the set difference removes the targeted elements together with the incident edges they sever. Random failure draws $R_{\text{rand}}(x)$ uniformly at random from $\mathscr{U}$ and reports Monte Carlo statistics across replicates; targeted failure selects $R_s(x)$ as the top-$K(x)$ elements under a centrality score $c_s : \mathscr{U} \to \mathbb{R}$:

$$R_s(x) = \text{TopK}\big(c_s, K(x)\big), \qquad s \in \mathscr{S} \setminus \{\text{rand}\}. \tag{3.34}$$

Ties in c_s are broken deterministically by element identifier so the entire sweep is reproducible.

Before reading the empirical results it is worth pinning down what the targeted failure does. The removable unit is structural, namely the consecutive-stop pair on the rail layer, the non-stop flight pair on the air layer, and the intra-FUA ground link on the coupling layer, which together form the universe $\mathscr{U}^{\text{edge}}$ of equation 3.32. Removing such an edge disables every scheduled service whose stop sequence traverses it, hence every supply edge those services carry through the map ρ of equation 3.37, hence every enumerated alternative that uses any of them. The structural removal universe, the supply cascade, and the service-aware score $\hat{b}$ of equation 3.36 therefore act on one indexed set of edges.

Centrality scores

The targeted removals rank edges by one of two centrality scores. The first is mode-blind and operates on the topological adjacency A^{topo} of equation 3.11,

$$c_{b^{\text{topo}}}(e) = \frac{1}{(|V|-1)(|V|-2)} \sum_{\substack{u \neq v \in V \\ u,v \notin e}} \frac{\sigma_{uv}(e)}{\sigma_{uv}}, \tag{3.35}$$

where σ_{uv} is the number of geodesics between u and v on A^{topo} and $\sigma_{uv}(e)$ the number passing through edge e (Brandes, 2001). Equation 3.35 is the standard unweighted Brandes edge betweenness on the consecutive-stop service graph of Section 3.1.3 and is denoted b^{topo} to disambiguate it from the service-aware score introduced below.

The service-aware betweenness takes account of all enumerated feasible itineraries. Let $\mathscr{B} = \{b = (\text{src}, \text{trg}, \text{dep})\}$ index the OD-and-departure buckets enumerated in Section 3.2.5, where $\text{dep} \in \{0, \ldots, 23\}$ is the hourly departure window. For each bucket b, denote by $p^{\star}_b$ the alternative attaining the minimum GTC and by $\Sigma(p^{\star}_b)$ the set of supply-layer edges it traverses. The service-aware edge betweenness $\hat{b}$ of an edge $e \in E^{\text{sup}}$ is

$$\hat{b}(e) = \big|\big\{ b \in \mathscr{B} : e \in \Sigma(p^{\star}_b) \big\}\big|, \tag{3.36}$$

namely the number of buckets whose lowest-GTC feasible itinerary runs through edge e. Equation 3.36 is the supply-side counterpart of the flow-weighted betweenness

of Cats et al. (2017). Their measure weights each link by the travel observed on it, which a study built from timetables rather than a demand matrix cannot do, so here each edge is weighted instead by the number of origin destination departure buckets that actually depend on it, and an edge scores highly when many lowest-cost itineraries are routed through it. Unlike b^{topo}, which counts shortest topological paths whether or not a service runs along them, $\hat{b}$ counts only edges that carry feasible scheduled itineraries under the timetable, calendar-overlap, and transfer-window constraints of Section 3.2.3.

Among the topological centralities, betweenness is preferred because vulnerability studies of public-transport networks find the largest connected component to be more sensitive to betweenness-based removal than to degree-based removal (Berche et al., 2009; von Ferber et al., 2009). Degree is local and misses the corridors that bridge otherwise separate regions, closeness varies little on a sparse spatial graph, and eigenvector centrality concentrates on dense cores that the European rail network lacks, whereas betweenness captures the shortest-path load whose loss propagates furthest. Plain topological betweenness is nonetheless an imperfect proxy for link importance in a scheduled system, because it treats every origin and destination as equally important and counts only shortest paths, so Cats and Jenelius (2014) recommend weighting betweenness by the service or passenger flow that actually traverses each link. The service-aware ordering $\hat{b}$ used here belongs to this flow-weighted family. It weights each edge by the number of origin destination departure buckets whose minimum-cost feasible itinerary traverses it, standing in for the passenger flows that a supply-side study does not observe. Comparing $\hat{b}$ against the topological ordering b^{topo} and treating their divergence as a result is how the analysis operationalises that recommendation.

Cascade on the supply layer.

For supply-layer scenarios, edge removal does not act directly on alternatives but on the services that traverse those edges. Let $\rho : E^{\text{sup}} \to 2^{\bigcup_\pi \mathcal{S}_\pi}$ map each supply edge to the set of services using it (equation 3.12), and let $\Pi_p \subseteq \bigcup_\pi \mathcal{S}_\pi$ denote the service set of alternative p. The disabled-service set under removal $R_s(x)$ and the surviving alternatives are

$$\mathcal{R}_{\text{dis}}(x) = \bigcup_{e \in R_s(x)} \rho(e), \qquad \mathcal{P}^{\text{surv}}(x) = \{ p : \Pi_p \cap \mathcal{R}_{\text{dis}}(x) = \emptyset \}. \tag{3.37}$$

The bucket-level effective GTC and its connected indicator are

$$\text{GTC}^{(x)}_{\min}(b) = \min_{p \in \mathcal{P}^{\text{surv}}(x) \cap b} \text{GTC}(p), \qquad \mathbf{1}^{\text{conn}}_b(x) = \mathbf{1}\big[\mathcal{P}^{\text{surv}}(x) \cap b \neq \emptyset\big], \tag{3.38}$$

with $\text{GTC}^{(x)}_{\min}(b) = +\infty$ when the bucket has no surviving alternative. Equations 3.37–3.38 make explicit that a single edge removal can invalidate many alternatives through the service-membership cascade, and that bucket survival is decided on the union of surviving alternatives rather than on a per-service basis.

Mode-dependent reroute classification

Once $\text{GTC}^{(x)}_{\min}(b)$ is known, every bucket b is assigned one of four states by comparing its post-disruption uplift $\delta_b(x) = \text{GTC}^{(x)}_{\min}(b) - \text{GTC}^{(0)}_{\min}(b)$ to a mode-dependent reroute

tolerance $T(b)$,

$$T(b) = \begin{cases} 60 \text{ min}, & \text{if } b \text{ is rail-only at baseline,} \\ 180 \text{ min}, & \text{if } b \text{ involves at least one air leg at baseline,} \end{cases} \tag{3.39}$$

and With baseline cost $g^{(0)}_{\min}(b)$, tolerance band $T(b)$, and unchanged tolerance ε, each bucket falls in exactly one state at intensity x:

$$\text{state}(b,x) = \begin{cases} \text{unchanged} & g^{(x)}_{\min}(b) \le g^{(0)}_{\min}(b) + \varepsilon, \\ \text{reroutable} & g^{(0)}_{\min}(b) + \varepsilon < g^{(x)}_{\min}(b) \le g^{(0)}_{\min}(b) + T(b), \\ \text{delayed} & g^{(0)}_{\min}(b) + T(b) < g^{(x)}_{\min}(b) < +\infty, \\ \text{disconnected} & g^{(x)}_{\min}(b) = +\infty. \end{cases} \tag{3.40}$$

with the unchanged tolerance set to $\varepsilon = 1$ minute. The two thresholds in equation 3.39 match the delay at which European law first entitles a passenger of each mode to compensation. For rail this is 60 minutes at the final destination, compensated under Regulation (EU) 2021/782 at one quarter of the fare and at one half beyond two hours (European Parliament and Council of the European Union, 2021). For air it is three hours, the delay the Court of Justice treats as equivalent to a cancellation under Regulation (EC) 261/2004 (Court of Justice of the European Union, 2009; European Parliament and Council of the European Union, 2004). The same 60 and 180 minute values are the generalised-travel-cost band widths of the path enumerator in Section 3.2.5, so a survivor whose uplift exceeds its band is treated as operationally disconnected. This refines the cut-off, delayed and unaffected passenger classification of Cats (2016) into four states by splitting its delayed class with a mode-specific reroute tolerance, so that 60 minutes remains meaningful for a rail leg without punishing an air leg the law allows to run three hours late.

Aggregate metrics

The bucket-level state in equation 3.40 is summarised at each edge removal fraction x by the disconnected and effective-disconnected fractions $f_{\text{disc}}(x)$ and $f_{\text{eff-disc}}(x)$, the average uplift on connected buckets $\Delta\bar{D}^{\text{conn}}(x)$, and the all-bucket average uplift with the disconnection penalty set equal to $T(b)$ per disconnected bucket,

$$\Delta\bar{D}(x) = \frac{1}{|\mathscr{B}|}\left[\sum_{b\,:\,\text{conn}} \min\big(\delta_b(x), T(b)\big) + \sum_{b\,:\,\text{disc}} T(b)\right]. \tag{3.41}$$

Equation 3.41 replaces the single global penalty of equation 3.26 by a per-bucket, mode-aware penalty. Bounding the per-bucket contribution at $T(b)$ ensures that the integral $Q = \int_0^1 \Delta\bar{D}(x)\,dx$ of equation 3.29 remains comparable across pools with different baseline GTC distributions, instead of being dominated by long disconnected transatlantic relations at high x.

The retained efficiency, normalised against the baseline so that $R(0) = 1$, is

$$R(x) = \frac{E_G(x)}{E_G(0)}, \qquad E_G(x) = \frac{1}{|\mathscr{B}|}\sum_{b\,:\,\text{conn}} \frac{1}{\text{GTC}^{(x)}_{\min}(b)}, \tag{3.42}$$

and the cumulative-fragility and accumulated-delay indices follow as in equations 3.28–3.29. Reporting $(F, Q, c_s^\star)$ together captures the area under the degradation curve, the total passenger-time penalty, and the location of the steepest drop, in the spirit of Cats et al. (2017).

Sweep grid and pool universes

Each scenario is evaluated on a uniform intensity grid $x \in \{0, \Delta x, 2\Delta x, \dots, 1\}$ with $\Delta x = 0.05$, giving 21 sweep points. A denser non-uniform grid concentrated near $x \in [0, 0.10]$ is used when the steepest portion of $R(x)$ falls in that region. Three pools are evaluated in parallel: a pure-rail pool over the 1,674 active intercity stations, a pure-air pool over the 197 active airports, and a unified intermodal pool over the union of rail-only, air-only, and rail–air alternatives produced by Section 3.2.5. The intermodal pool is the operational instantiation of the supply layer of equation 3.13 restricted to the bucket set $\mathscr{B}$ and is referred to as the intermodal pool in the rest of the chapter. A bucket survives in the intermodal pool whenever any of its rail-only, air-only, or mixed alternatives survives the cascade of equation 3.37, so that intermodal robustness is bounded below by single-mode robustness by construction.

3.4. Scenario taxonomy

The study aims to provide policy implications relevant to future European network development. Accordingly, it is not a pure complex-network-science exercise but rather a comprehensive analysis that respects the operational realities of traffic; each removal must correspond to a recognisable mechanism in the real system (M. Zhang et al., 2022). The scenarios are classified into three types, following the principle that conventional random and targeted failure scenarios serve as benchmarks while mechanism-based scenarios are designed to reflect plausible real-world disruptions. Type I applies structural stress tests on the supply layer, Type II models institutional disruption through operator withdrawal and cooperation-tier restriction, and Type III maps a documented macro-hazard into the set of services it disables.

Edge removal, within the scenario taxonomy of this study, acts as a precise intervention on the capacity or availability of services, often over a time window. On the service network, edge removals correspond to the cancellation of the associated scheduled departures. Node or station failures are not represented as abstract deletions, but as bundled removals of the incident service and transfer edges associated with the affected facility. This taxonomy ensures that edge-removal experiments remain interpretable in operational terms and comparable across network configurations.

3.4.1. Type I: Structural stress tests

Structural stress tests serve as reference benchmarks against which mechanism-based scenarios are evaluated. They include two sub-types.

Random failure removes nodes or edges uniformly at random, simulating the baseline failure rate of a network with no preferential targeting. The random scenario is repeated across multiple independent realisations to obtain confidence intervals, since individual random removals can produce widely varying outcomes depending on which elements happen to be selected.

Importance-based failure removes nodes or edges in descending order of a centrality metric. Two edge orderings are applied on the supply layer: the itinerary-anchored service-aware betweenness $\hat{b}(e)$, which counts the origin-destination-departure buckets whose fastest-cost alternative traverses edge e, and the unweighted topology betweenness $b^{\text{topo}}(e)$. The divergence between these two orderings is itself a central result of the analysis. Importance-based removal represents a worst-case stress test, revealing the maximum damage that a targeted disruption could inflict on the network. The comparison between random and importance-based results directly tests the robust-yet-fragile hypothesis from scale-free network theory (Barabási, 2016).

3.4.2. Type II: Institutional disruption through operator withdrawal and cooperation tiers

Type II scenarios act on the operators that provide supply-layer services rather than on individual edges, and capture the institutional dimension of European long-distance robustness. Two variants are evaluated. Operator withdrawal removes every supply-layer edge served by a designated operator subset, treating the full set of services an operator runs as a single failure unit, and is reported for single-operator ($k = 1$), pairwise ($k = 2$), and triplet ($k = 3$) withdrawals across the rail train operating companies and air carriers in the network and this models bankruptcy, strike action, or the regulatory withdrawal of an operator's licence.

Where operator withdrawal removes a chosen set of operators, the cooperation-tier experiment is defined by which operators a passenger may use. An itinerary is retained only when every leg is run by a member of a defined cooperation set C, equivalently the routes of every operator outside C are withdrawn. The set C is widened through five arrangements drawn from current European cooperation practice as characterised by Beria et al. (2023). The ladder runs from a single carrier's ticket (B1-strict), the limited cross-operator entitlement that European Parliament and Council of the European Union (2021) grants today, through one permitted change of carrier (B1-split), the documented bilateral through-tickets such as the Deutsche Bahn and ÖBB Nightjet (B2), and the Railteam alliance of European high-speed operators (B3) (Railteam, 2026), to full rail interlining as the continental analogue of the airline multilateral interline agreements (B4) (International Air Transport Association, 2026) and full rail and air interlining (B5), the integrated booking that Cats (2025) identify as still unrealised and which sets the no-disruption baseline. Each tier is scored by retained efficiency and the four-state outcome on the surviving itineraries, so the B1-to-B5 progression isolates the robustness conferred by the passenger's rebooking entitlement rather than by the infrastructure, with C treated as a first-class parameter varied alongside the disruption itself.

3.4.3. Type III: Macro-scenario disruption

Type III scenarios translate a documented macro-hazard into the set of services it disables, and test whether structural redundancy or cross-modal substitution restores robustness under a real spatially correlated shock.

Two scenario families are reported. A continental-airspace closure modelled on the April 2010 Eyjafjallajökull eruption disables every flight with an endpoint inside the day-resolved ash-zone polygons, testing rail as a substitute for air. A continental-flood sce-

nario, modelled both as a river-basin-polygon sweep and as the Copernicus rapid-mapping of the July 2021 Bernd storm, disables the rail and air services whose facilities fall inside the inundated area, testing air as a substitute for rail.

The ash cloud covers almost the whole air layer at once, the macro-shock counterpart of random failure, while the flood concentrates on one small region of the dense rail network in Western Germany and Benelux . The pair therefore tests the two ways a real hazard can fall on the network, spread across the whole of it or clustered on one part.

4

The European multilayer transport network

This chapter describes the data used by the robustness analysis of Chapter 5, and it proceeds in four steps. Section 4.1 fixes the operationally active rail and air layers and describes their operator and carrier structure. Section 4.2 characterises the topology of those layers and compares it with the prior literature. Section 4.3 reports the path enumeration that produced the feasible itinerary pool. Section 4.4 validates the enumerated travel times against an independent reference. The chapter is descriptive throughout. Where the topology suggests how the network will behave under disruption, that reading is stated as a prediction to be tested later rather than as a result.

4.1. The operational network over the study area

Section 3.1 established the multilayer representation of the network, comprising the structural multilayer network of Section 3.1.3 and the functional multilayer network of Section 3.1.4, and named its two inputs, the GTFS feeds of the long-distance rail services and the commercial flight schedules. This section describes what those inputs contain once they are reduced to the operationally active sub-network that the robustness analysis uses. It reports the size of each layer, the countries each one covers, the data coverage quality by country, and the operators that run the services, across the 30 countries of the study area defined in Section 1.3.

4.1.1. Network scale and spatial coverage

Each of the two layers is reduced to its operationally active subset before any further analysis is performed, through a filter applied within the 30 countries of the analysis list defined in Section 1.3 and Table 1.1. A rail station is retained only when it is traversed by at least one rail-mode route and appears in at least one rail-mode trip in the timetable, a criterion that retains 1,674 active stations across 26 countries. An airport is retained when it is geographically located within or near an FUA and carries at least one scheduled

departure in the air timetable, a filter that retains 197 operating airports across 30 countries. Each station is assigned to the country of its FUA, and each airport to the country in which it lies. An airport can serve more than one FUA, so it is counted once under the country that contains it rather than under every area it serves. The two counts, 1,674 stations and 197 airports, are the node set carried through every later table and figure.

These stations and airports, together with the edges between them, are the empirical realisation of the multilayer formalism of Section 3.1. The stations form the rail node set V^r and the airports the air node set V^a of the structural multilayer network G^{s}, whose edge set E^{s} collects the consecutive-station rail edges, the airport-pair air edges, and the ground-access edges that couple the two modes within a functional urban area. The same node set carries the functional multilayer network G^{f}, whose topological layer is the structural supra-adjacency $\mathbf{A}^{\mathrm{s}}$ of equation 3.11 and whose supply layer E^{sup} holds the scheduled services. The global topology of the structural layers is characterised in Section 4.2, and the supply layer feeds the enumerated itinerary pool of Section 4.3. Treating the empirical network as the instance of that formal object is what lets a topological reading and a service-level reading of the same network be compared rather than treated as separate studies.

Table 4.1 reports the per-country node and route counts together with structural indicators for both the rail and the air layer, and Figure 4.1 shows the same coverage graphically. Both layers are unevenly spread, but their concentrations sit in different places. On the rail layer Germany alone accounts for 39% of all stations and 41% of all routes, followed by United Kingdom at 14% of stations and France at 8%, while nine countries each hold fewer than ten stations and Romania and Finland carry no cross-border edges at all. On the air layer the leaders are different and the spread is gentler. Italy, United Kingdom and Spain hold the most airports at 27, 26 and 25, Germany and France follow at 21 and 20, and these five countries hold 119 of the 197 airports, a concentration far milder than the single-country dominance that Germany exerts on rail. The air layer also reaches where the rail layer is thin, with Spain served by 25 airports against 48 stations and with Bulgaria, Cyprus, Greece and Malta entering the network through the air layer alone.

Three structural patterns follow from Table 4.1 and are made visible in Figure 4.2. First, Germany and the United Kingdom combine the highest station counts with a high supply degree, 33.7 and 30.0 against a mean of 25.6, so the average station reaches that many others without a transfer. This dense direct connectivity provides many alternative routings, so most pairs stay connected when a service is disrupted, and Belgium and the Netherlands show the same pattern on a smaller station base. France is the exception, and the contrast between its two degrees explains it. Its stations carry the highest structural degree, 5.9, so the typical station has about six others as consecutive stops, yet its supply degree of 15.8 is the lowest of the five largest networks. French long-distance rail is dominated by high-speed services that link major cities directly and stop at few stations in between. A station on several such services has many distinct consecutive stops, which raises its structural degree, but each service makes so few stops that the number of stations reached without a transfer stays small. Germany and Austria do the reverse, with services that stop often enough to put 33 to 34 stations within a no-transfer ride of the typical station.

Second, the cross-border route share rises as the scale of national network shrinks. The smallest networks are almost entirely international, with Luxembourg, Croatia, Latvia

and Estonia at 100% and Slovenia and Slovakia at 86% and 79%, so nearly every route based in those countries crosses a border. The rail networks of Austria and Czechia have moderately high cross-border shares, at 45% and 55% respectively, whereas the largest networks are overwhelmingly domestic, with Germany and France both below 10%. Latvia and Estonia are reduced to a single active station each and enter the network only through their international corridors. Third, several peripheral countries such as Romania, Finland and Portugal combine low degree in both layers with minimal cross-border connectivity, whether through geography or through the absence of neighbouring-country data, which leaves them structurally vulnerable. Ireland is likewise separated from the continental networks, although it runs a cross-border line into Northern Ireland. The air panel of Figure 4.2 shows a complementary structure on the other layer. A few countries with only one or two operating airports, such as Switzerland, Portugal and the Netherlands, nonetheless reach very high routes-per-airport, which is the hub-and-spoke signature that the air operator structure of the next subsection makes explicit, while the large multi-airport networks of Italy, United Kingdom, Spain, Germany and France spread their routes across many airports at a lower density.

4.1.2. The operators and their geographic reach

The supply layer E^{sup} of the functional multilayer network carries the scheduled services, run by two sets of operators, the rail train operating companies and the air carriers. Once the ground-access links between stations and airports are added, the two sets can be considered as a whole. Each layer is dominated by a few operators, as Figure 4.3 shows, but where their routes go and how far they reach differ sharply and in opposite directions. For the rail networks, the assembled GTFS feed identifies 51 train operating companies that run 30,451 distinct routes, of which 28,248 stay within one country and 2,203 cross a border, a domestic share of 92.8%. One operator stands out. DB runs 11,636 routes, or 38.2% of the total, and only 531 of them are international, an international share of 4.6%. SNCF runs 3,895 routes or 12.8%, DFTO 3,311 or 10.9% all within Great Britain, and Trenitalia 1,305 or 4.3%, so these four operators run 66.2% of all routes. A small group of incumbents carries the cross-border services. ÖBB reaches 14 countries and DB 13, the two widest reaches. Eurostar runs 122 international routes against a single domestic one, and ZSSK and CFL run mostly across borders, at international shares of 89.0% and 97.8%, in line with the corridor character of their home networks.

The air layer carries a structure that is concentrated in the same way but pointed the other way. The March 2025 schedule is flown by 31 carriers over 10,203 directed routes, of which only 843 are domestic and 9,360 are international, an international share of 91.7%. The low-cost carriers dominate. Ryanair flies 3,592 routes or 35.2% of the total, easyJet 1,342 or 13.2%, Eurowings 449 or 4.4% and Lufthansa 394 or 3.9%, so the four largest carriers fly 56.6% of all routes. The air network is a hub-and-spoke structure built on the 197 operating airports, of which 51 are large hubs and the remainder are medium or small, and the long edges of Figure 4.4b radiate from a small number of those hubs. The contrast with rail is the central point of this subsection. Rail supply is concentrated in national incumbents that run mostly within their own borders at a domestic share of 92.8%, whereas air supply is concentrated in pan-European low-cost carriers that run mostly across borders at an international share of 91.7%. The two layers therefore present a traveller with

Table 4.1: Rail and air network structural statistics by country, sorted by descending rail-station count. The 26 countries with an active rail layer are followed by four countries that enter the robustness network only through the air layer.

		Rail					Air			
ICC	Country	Stations	Routes	$\bar{k}_{str}$	$\bar{k}_{sup}$	CB%	Airports	Routes	$\bar{k}$	Intl%
DE	Germany	654	13,235	5.3	33.7	8.3	21	837	26.5	90
UK	United Kingdom	241	7,362	5.1	30.0	0.1	26	1,317	29.4	91
FR	France	140	3,857	5.9	15.8	7.4	20	814	27.2	77
IT	Italy	118	1,497	4.8	18.8	7.7	27	1,180	33.7	74
PL	Poland	80	739	4.7	19.3	8.5	12	440	27.5	93
CH	Switzerland	78	775	5.0	19.0	7.2	2	239	67.5	99
BE	Belgium	70	660	3.9	26.3	23.6	2	184	60.5	99
NL	Netherlands	61	697	5.0	24.1	9.8	4	270	39.5	100
ES	Spain	48	363	3.5	10.1	0.8	25	1,071	29.4	83
HU	Hungary	35	241	3.4	11.7	5.4	1	120	80.0	100
RO	Romania	35	102	2.7	7.4	0.0	6	80	11.8	87
AT	Austria	22	406	5.2	33.9	44.6	6	241	23.2	97
SE	Sweden	18	191	4.9	10.8	30.9	5	163	20.4	94
DK	Denmark	12	59	2.8	10.1	3.4	2	27	11.0	100
CZ	Czechia	11	38	4.5	15.0	55.3	3	116	25.7	99
PT	Portugal	11	38	2.5	6.1	2.6	2	241	65.5	99
NO	Norway	10	27	2.6	4.4	25.9	5	192	19.2	83
FI	Finland	8	42	2.2	4.5	0.0	5	91	11.2	92
SK	Slovakia	5	14	5.2	15.6	78.6	3	29	8.7	96
LU	Luxembourg	5	20	4.6	20.0	100	1	66	44.0	100
IE	Ireland	4	8	1.8	1.8	25.0	3	165	44.0	99
HR	Croatia	2	5	2.5	7.5	100	3	73	19.3	96
LT	Lithuania	2	35	1.5	1.5	2.9	2	71	27.5	100
SI	Slovenia	2	29	3.5	8.5	86.2	1	10	10.0	100
LV	Latvia	1	2	2.0	2.0	100	1	77	58.0	100
EE	Estonia	1	1	1.0	1.0	100	1	35	28.0	100
air only, no active rail layer										
BG	Bulgaria	0	0	–	–	–	3	104	24.7	96
CY	Cyprus	0	0	–	–	–	2	88	34.5	100
EL	Greece	0	0	–	–	–	2	190	52.0	99
MT	Malta	0	0	–	–	–	1	100	74.0	100
	Total	**1,674**	**30,451**				**197**	**8,631**		

$\bar{k}_{str}$ is the mean degree in the structural layer, where two stations are adjacent when they are consecutive stops on a service, and $\bar{k}_{sup}$ is the mean degree in the supply layer, where two stations are adjacent when one service serves both. For air every flight is non-stop, so the structural and supply layers coincide and the single $\bar{k}$ reported is the airport degree. Each rail route is counted once, under the country where most of its stops lie, so the Routes column sums to the distinct route total rather than double-counting cross-border routes. The 26 country rows sum to 30,443, and the remaining 8 of the 30,451 distinct routes lie entirely in micro-states or the four air-only countries and are not shown in a row. For rail, CB% is the cross-border route share, the percentage of a country's routes that cross a national border, so a high CB% marks a network that works mainly as a corridor between neighbours and a low CB% one that works mainly within its own borders. For air, Intl% is the international route share, and the air routes are directed routes by country of origin, the 8,631 originating in these 30 countries being the study-area part of the 10,203 routes in the full schedule. A dash in the rail block marks a country with no active rail layer. The 197 operating airports are the 189 in the 26 rail countries together with the 8 in the four air-only countries.

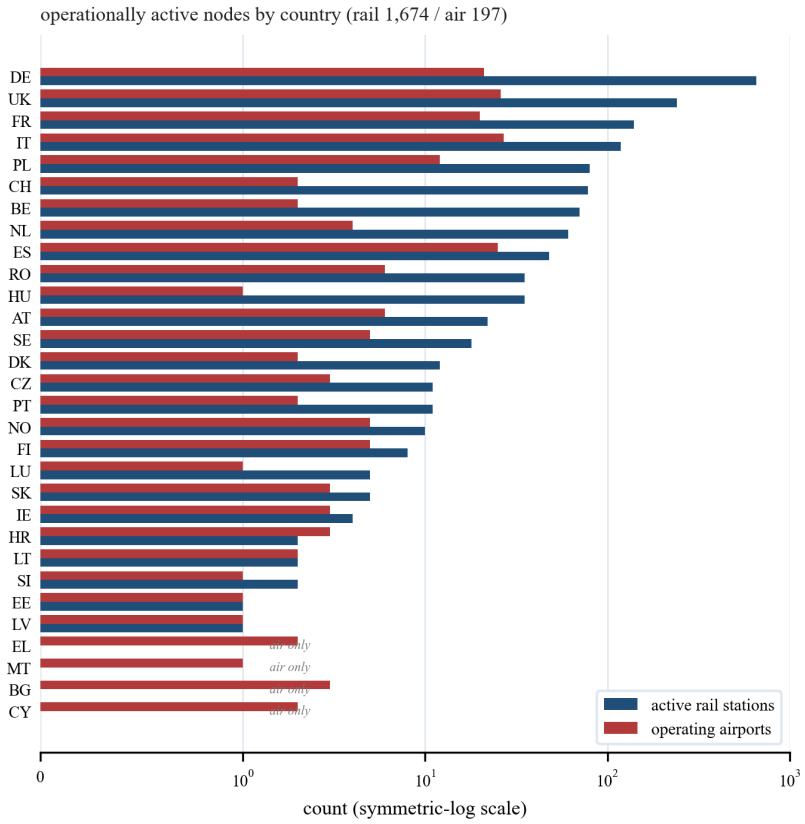

Figure 4.1: Operationally active nodes by country on a symmetric-log scale, with active rail stations in navy and operating airports in red. The rail layer is concentrated in a small number of large western networks while the air layer is comparatively flat across countries, which is the first sign of the asymmetry that shapes the enumerated pool.

opposite geographies, a dense national rail layer for short and medium distances and a sparse international air layer for long ones, and the enumerated pool of Section 4.3 is the product of that complementarity.

The counts above are route definitions and describe the supply at the level of the timetable. The robustness analysis of Chapter 5 acts on a different and complementary object, the set of supply edges each operator serves on the intermodal pool, and it treats the air carriers on the same terms as the rail operators. That analysis works with 82 operators in total, 51 rail operators and 31 air carriers, namely the 51 train operating companies identified above. On that supply-edge basis DB again stands apart, serving 1,499 supply edges of the intermodal pool, with the low-cost carrier Ryanair second at 1,364 edges and the remaining operators an order of magnitude below. The descriptive route count of 11,636 for DB and the supply-edge count of 1,499 therefore measure two different things, the number of scheduled route definitions and the number of distinct station-pair edges

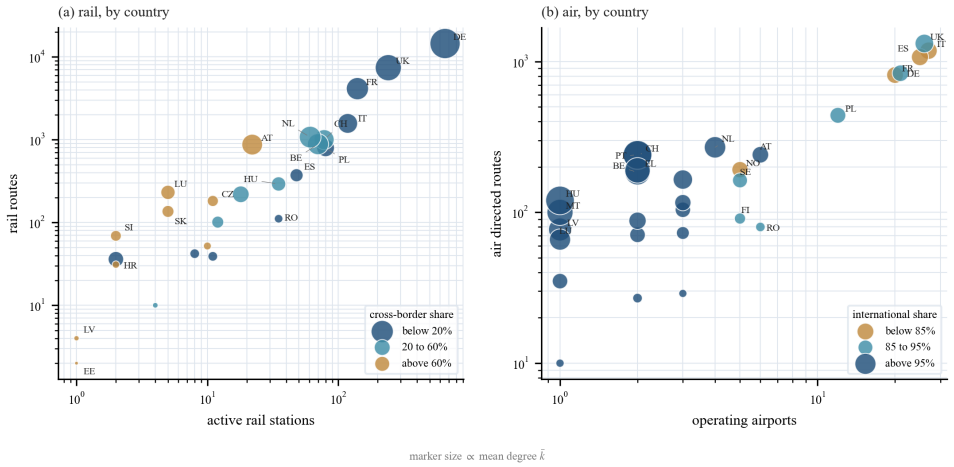

Figure 4.2: Network size and density by country on log axes, with marker size proportional to mean degree. Panel (a) shows the rail layer as active stations against routes, coloured by cross-border route share. Panel (b) shows the air layer as operating airports against directed routes, coloured by international route share. Germany dominates the rail panel by an order of magnitude, while the air panel separates the large multi-airport networks of Italy, United Kingdom, Spain, Germany and France from the small single-hub-and-spoke networks such as Switzerland, Portugal and the Netherlands that reach very high routes-per-airport.

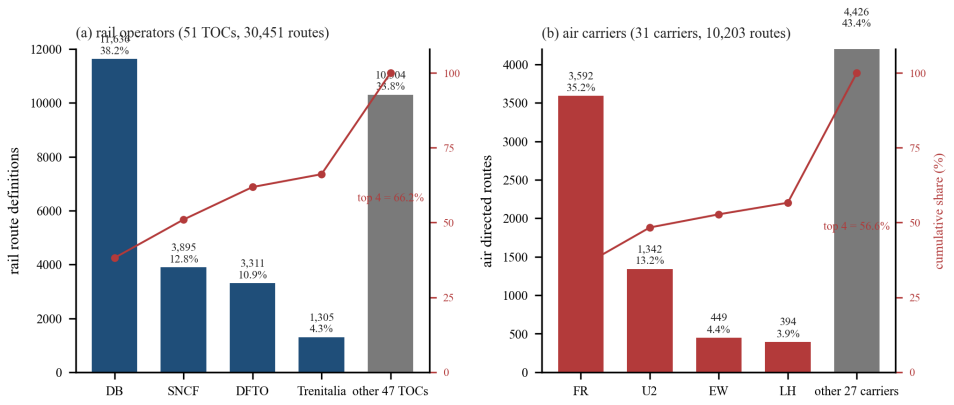

Figure 4.3: Route concentration on each layer, with the cumulative share overlaid in red. Panel (a) shows the four largest rail operators against the remaining 47 on the route-definition basis, where the four largest define two thirds of all routes and DB alone two fifths. Panel (b) shows the four largest air carriers against the remaining 27 on the directed-route basis of the March 2025 schedule, where the four largest fly 56.6% and Ryanair alone 35.2%.

that survive into the feasible pool, and both point to the same concentration.

As part of the calibration after data collection, the observed operator structure was compared with the 69 long-distance city pairs analysed by Beria et al. (2023). The comparison agrees with that study in placing the most developed pre-fourth-package open-access cases in Italy and Czechia, with Italy the exceptional case of a full high-speed challenge and Czechia characterised by multi-operator competition between ČD, RegioJet and Leo Express. A corridor duopoly appears in Austria on Vienna to Salzburg between WESTbahn and ÖBB, and Sweden appears as an early mature open-access case centred on Stockholm to Gothenburg. Germany, despite 22 active operators, remains dominated by DB, whose 11,636 routes dwarf the 153 of Flixtrain. The feed does not reproduce the market sample of the literature exactly, but it reflects the same uneven geography of liberalisation in which richer operator overlap is concentrated in a limited number of corridors and countries. That concentration is expected to matter for robustness, because a corridor served by several operators carries built-in service redundancy while a single-operator corridor loses all service when that operator withdraws. The Great Britain sub-network is the clearest case. The feed shows 11 domiciled operators and 13 active operators, in line with the characterisation of Britain by Smith and Nash (2023) as the most fragmented rail system in Europe. The prominence of DFTO with 3,311 domestic routes reflects the post-pandemic arrangements under which the government took control of failing franchises, so its dominance is a transitional artefact of the crisis period rather than a stable market feature.

4.2. Network characteristics

The robustness signature of any network is conditioned on its degree distribution. Albert et al. (2000) showed that scale-free graphs collapse under the targeted removal of their highest-degree stations yet tolerate random failure, whereas exponential graphs degrade in nearly the same way under either regime. A robustness curve is therefore uninterpretable until the degree distribution of each layer is fixed. Broido and Clauset (2019) further showed that strict scale-freeness over the full degree range is rare in real networks, and that the usual empirical case is a heavy tail with an exponential cutoff. The aim here is not to declare each layer scale-free or not but to place each layer on that spectrum using the maximum-likelihood tail-fitting protocol of Clauset et al. (2009) rather than a biased log-log regression. The result is a topological prior that the robustness chapter then tests.

Figure 4.4 maps the two layers and their union. The rail layer is a continental web of short consecutive-station edges clustered into national networks, while the air layer is a sparse set of long edges radiating from a small number of stations, and the union inherits both. The visual contrast between the dense clustered rail layer and the sparse long-range air layer is the geographic counterpart of the topological indicators that follow.

The characterisation works with four graphs drawn from the formalism of Section 3.1. The rail layer, written rail (L) and following the infrastructure-space convention of von Ferber et al. (2009), is the rail block G^r of the structural multilayer network, with the 1,674 stations as nodes and an edge between any two stations that are consecutive on a service. The air layer is the air block G^a, where every flight is non-stop so the infrastructure space and the service space coincide. The intermodal layer is the structural supra-adjacency $\mathbf{A}^s$ of equation 3.9, the union of the rail and air edges with the within-area cross-mode edges, which is by construction the topological layer of the functional multilayer network.

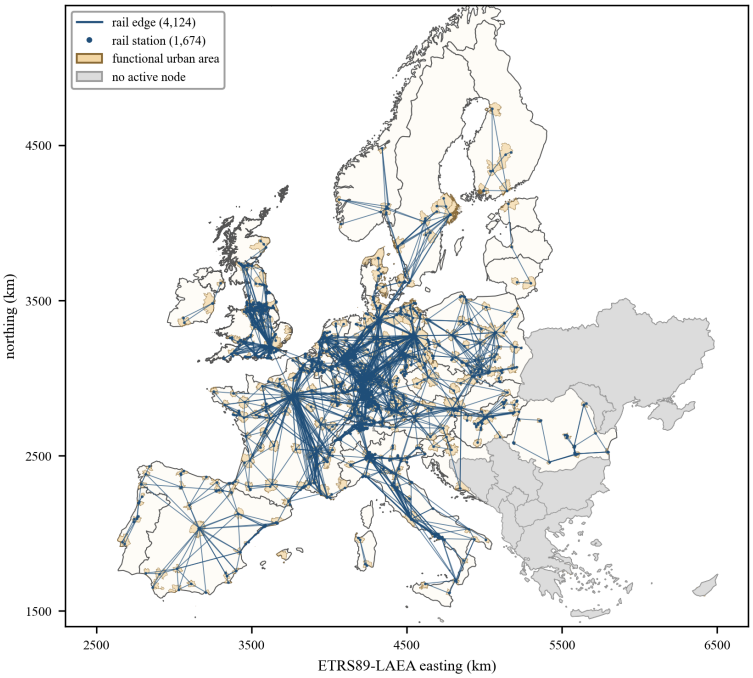

(a) rail layer

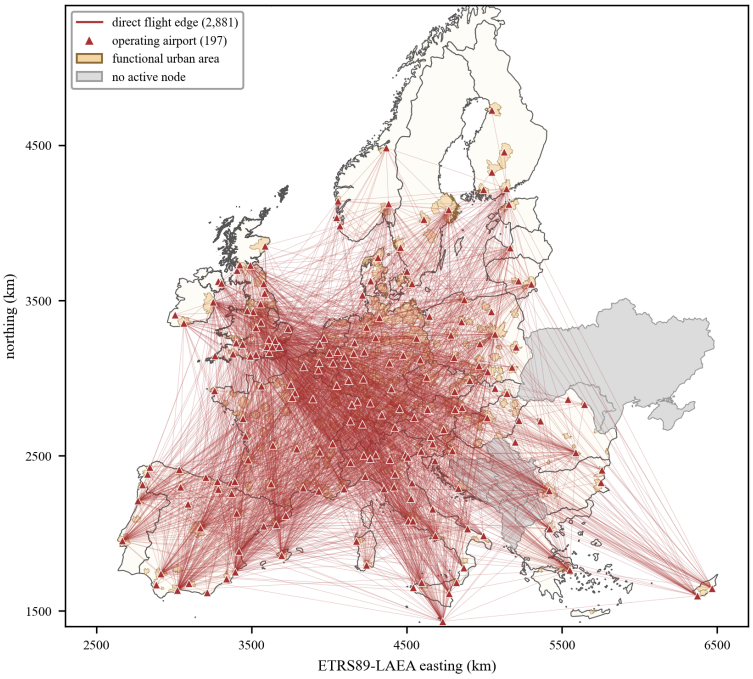

(b) air layer

Figure 4.4: The two layers of the operational network over the study area, with the countries that carry no active node of a layer shaded in grey in that layer's panel. Panel (a) shows the 1,674 rail stations and their 4,108 consecutive-station edges, clustered into national networks. Panel (b) shows the 197 operating airports and their 2,881 direct-flight edges reaching across the whole continent.

The fourth graph, written rail (P) and following the service-space convention, connects any two stations served by a common route, and it is the simple-graph projection of the supply-layer relation of Section 3.1.4, namely two terminals reachable without alighting on the same service. Reading rail (L) against rail (P) on the same 1,674 stations therefore measures how much the supply layer adds to the bare topological layer, which is the empirical version of the structural-to-functional distinction that the methodology draws.

Table 4.2 reports the global indicators on each layer. The four numbers that drive the reading are the heterogeneity ratio $\kappa = \langle k^2 \rangle / \langle k \rangle^2$, the assortativity r, the diameter d, and the gap between the rail infrastructure space and the rail service space. The rail infrastructure space combines a low κ of 2.16, positive assortativity of +0.26, and a long diameter of 22. The air layer combines a low κ of 1.90 with a much heavier-than-random tail, a maximum degree of 108 on only 230 graph nodes, strong disassortativity at −0.22, and a short diameter of 4. The intermodal layer reaches $\kappa = 3.92$, heavier than either single layer, with a maximum degree of 192 located where a major airport and a major station fuse in the same urban area, as at Frankfurt, Paris, Amsterdam, London and Madrid. The comparison between the rail infrastructure space and the rail service space on the same stations multiplies the edge count fivefold, raises clustering from 0.40 to 0.72, and shortens the average path length from 6.5 to 3.8. That gap measures the connectivity the service overlay supplies on top of the bare track.

Table 4.2: Global indicators for each layer. All values use the full graph except clustering, assortativity, average path length and diameter, which use the largest connected component. The air and intermodal layers carry a small number of isolated graph nodes, so the largest connected component holds 197 of the 230 air graph nodes.

indicator	rail (L)	rail (P)	air	intermodal
$\lvert V \rvert$	1,674	1,674	197	1,907
$\lvert E \rvert$	4,108	21,419	2,881	10,013
LCC fraction	96.4%	96.4%	85.7%	99.3%
$\langle k \rangle$	4.91	25.56	29.25	10.57
$k_{\max}$	53	256	108	192
$\kappa = \langle k^2 \rangle / \langle k \rangle^2$	2.16	2.45	1.90	3.92
$\langle C \rangle$	0.400	0.721	0.560	0.520
assortativity r	+0.26	+0.27	−0.22	+0.07
$\langle \ell \rangle$	6.53	3.83	2.04	3.52
diameter d	22	11	4	11
γ-connectivity	0.82	4.26	4.21	1.75
α-meshedness	0.73	5.90	5.83	2.13
mean closeness	0.160	0.270	0.502	0.291
mean directness	1.19	1.14	1.72	1.09

4.2.1. Degree and betweenness distributions

Table 4.3 reports the maximum-likelihood power-law fits and the log-likelihood ratios against three alternatives, and Figure 4.5 shows the four degree distributions with the fitted tail overlaid. A negative ratio with a small p-value means the alternative fits the tail strictly better than a pure power law, while a near-zero ratio with a large p-value means

the two candidates cannot be distinguished and the tail cannot be called scale-free. The rail infrastructure space is statistically indistinguishable from an exponential distribution, with a ratio of +0.55 at $p = 0.58$. The rail service space is likewise indistinguishable from an exponential, with a ratio of +0.48 at $p = 0.63$, so the service overlay produces small-world reachability without producing a scale-free distribution. The air layer carries the strongest hub structure of the four, with a fitted exponent of 2.98 in the scale-free range $2 < \hat{\alpha} < 3$ (Albert et al., 2000; Barabási, 2016), a maximum degree of 108 against a mean of 29, and disassortative mixing that concentrates connectivity on the legacy hubs. Its degree tail is nonetheless not a pure power law, since both the exponential (ratio −6.01 at $p < 10^{-8}$) and the lognormal (ratio −3.34 at $p < 10^{-3}$) fit it strictly better, the signature of an exponential high-degree cutoff on the hub-and-spoke geometry. The intermodal layer decisively beats the exponential but cannot be separated from a lognormal, which is exactly the heavy-tailed-with-cutoff regime that Broido and Clauset (2019) identify as the empirical norm.

Table 4.3: Maximum-likelihood power-law tail fits (Clauset et al., 2009). D_{KS} is the Kolmogorov-Smirnov distance, LR is the normalised log-likelihood ratio against the named alternative, p is its two-sided significance, and n_{tail} is the number of nodes at or above $x_{\min}$.

layer	$\hat{\alpha}$	$x_{\min}$	D_{KS}	n_{tail}	LR vs alternative (p)
rail (L)	2.28	3	0.063	1,029	exp +0.55 (0.58), logn −5.36 ($< 10^{-7}$)
rail (P)	2.75	34	0.049	380	exp +0.48 (0.63), logn −2.47 (0.014)
air	2.98	37	0.189	68	exp −6.01 ($< 10^{-8}$), logn −3.34 ($< 10^{-3}$)
intermodal	2.28	6	0.021	929	exp +9.52 ($< 10^{-20}$), logn −1.07 (0.29)

Degree is a property of nodes, whereas the robustness experiments of Chapter 5 remove edges. The two connect through edge betweenness, since the edges incident to a high-degree node carry a disproportionate share of it, so the degree-distribution class of each layer is the node-level precursor of the topological edge-betweenness ordering b^{topo} that the results chapter uses as one of its two removal sequences. Whether that degree-based prior survives is itself a result, because the results chapter finds that the topological ordering behaves almost like random edge removal while only the service-aware ordering $\hat{b}$ inflicts disproportionate damage. The degree distribution therefore sets the expectation against which the service-level reading is measured rather than a conclusion in its own right.

The betweenness distribution shows that air and the rail infrastructure space share a

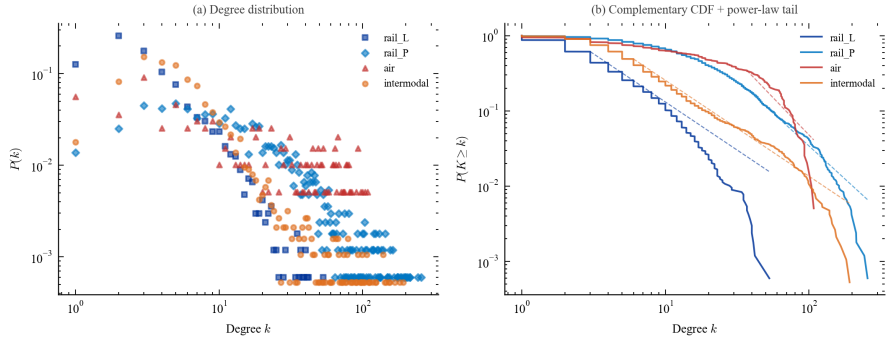

Figure 4.5: Degree distribution and complementary cumulative distribution for the four layers, with the fitted power-law tail overlaid as a dashed line on each complementary cumulative distribution. Only the air layer produces a tail that beats every alternative model at high significance.

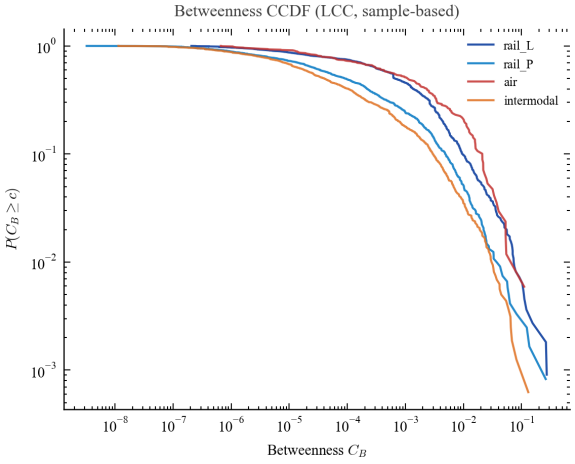

Figure 4.6: Complementary cumulative distribution of betweenness centrality on the largest connected component. Air and the rail infrastructure space share a heavy betweenness tail despite very different degree distributions, because in both layers a small set of stations carries a large share of shortest paths.

heavy betweenness tail despite very different degree distributions, because in both layers a small set of stations carries a large share of shortest paths. This is why Šfiligoj et al. (2026) treat heterogeneity in betweenness as the signature of a hub-and-spoke structure. For transport networks betweenness is often more discriminating than degree, because a station can hold a low degree yet sit on many shortest paths, with Lyon, Lille and Brussels the textbook European high-speed examples.

4.2.2. Comparison with the prior literature

The indicators above can be read against the topology-of-transport-networks literature, where they agree with some findings, sharpen others, and sit in tension with a few. Reading them this way serves two purposes. It places the present network within an established

body of empirical work, and it converts the descriptive indicators into a set of expectations that the robustness analysis of Chapter 5 can later confirm or overturn. Five comparisons matter, and they move from the rail layer through the air layer to their combination.

The exponential degree tail of the rail layer places it with the recent empirical high-speed-rail studies rather than with the early and now-doubted claim that transport networks are generically scale-free. Li and Rong (2020) report a Chinese high-speed degree distribution that is close to uniform on 32 stations and 340 service edges and never resolves into a clean power law, and Jiao et al. (2020) model the same system in service space and find a similarly truncated structure that is most robust under random removal and least robust under betweenness-based removal. The European rail layer behaves the same way, and the broader argument of Broido and Clauset (2019) that strict scale-freeness is rare in real networks is borne out by the European input and the prior Chinese studies alike. The consequence to carry forward is that the rail layer should not be expected to show the signature scale-free collapse under degree-targeted removal, simply because it has no degree hub for such a removal to find.

The air layer is the opposite case, and it is the one place where the scale-free description genuinely holds. Its fitted exponent of 2.98 and the decisive rejection of the exponential alternative sit squarely in the range that Pien et al. (2015) report for the European air-traffic network. Those authors decline to fit a formal power law and caution that the highest-betweenness nodes are not always the busiest in throughput, because the control structure of European airspace partly decouples topological centrality from realised flow (Guimerà et al., 2005). That caution does not contradict the scale-free signature found here. It refines how the signature should be used, as a structural property of the graph rather than as a promise that the highest-degree airports are the most disruption-critical, and it is the reason the robustness chapter orders the failures by both a degree-based and a flow-based centrality rather than by degree alone.

What that concentration implies for robustness has a close precedent in the air sector itself. Morlotti and Redondi (2021) test the four largest European air-freight integrators and find that the most concentrated hub-and-spoke carrier loses 92% of its connectivity when its principal station is removed and is the least robust of the four, while the most point-to-point carrier is the most robust. The mechanism, that a heavy concentration of traversal at a few stations turns into a steep loss once those stations are removed, is the same mechanism that the heavy air betweenness tail of Figure 4.6 sets up here, even though the passenger air heterogeneity ratio is below two. The expectation that follows, before any robustness curve is computed, is that the passenger air layer will be acutely exposed to the removal of its busiest few stations.

Where the rail layer is exposed is a separate question with a separate answer, and here the present indicators agree with both the modelling and the historical literature. Jiao et al. (2020) report that a disruption in the densest region of the Chinese network produces the largest loss of weighted efficiency and that a single long-corridor closure causes a drop of 68%. A long diameter, a positive assortativity and a low planar-connectivity index are precisely the conditions under which a graph is fragile to spatially correlated removal rather than to random degree removal, and the European rail layer carries all three. This gives a quantitative counterpart to the qualitative argument of Toet et al. (2026) that European rail grew as a collection of national networks focused on domestic travel, an argument that

the operator structure of the previous section already supported from the supply side.

The final comparison concerns what happens when the two layers are read as one, and it is the comparison most directly tied to the research question. Ippolito and Cats (2024) build a comparable rail-and-air multilayer graph for Europe and find a high area-under-curve score for the rail layer under random failure, a much lower score under degree-targeted failure on the air layer, and the lowest score of all for the combined graph under betweenness-targeted failure. The intermodal indicators reported here explain both halves of that result. The combined heterogeneity ratio of 3.92 and maximum degree of 192 are heavier than either single layer, which steepens any targeted-failure curve, while the rise in connected fraction to 99.3% against 96.4% for rail and 85.7% for air explains the high score under random failure. The expectation carried into Chapter 5 is therefore twofold. The rail-air combination should buy robustness against random and mode-specific shocks, and it should at the same time concentrate a new exposure at the small set of cross-mode stations where the two layers fuse, so that a betweenness-based failure aimed at those stations comes to dominate the intermodal vulnerability story.

4.2.3. Topology-based disruption hypotheses

The findings above fix the prior under which the robustness chapter reads its disruption curves, summarised in Table 4.4. The rail infrastructure space should be most affected by single-edge failures and by geographically correlated removals, because a long diameter, a low γ index and a positive assortativity leave it exposed to single-corridor closures. The air layer should be most affected by degree-targeted and high-betweenness failures, because it is the only layer with a genuinely scale-free tail. The intermodal layer should rebalance both, because its heavier combined tail and its cross-mode edges raise the connected fraction to 99.3% while concentrating traversal at a small set of cross-mode stations. The comparison between the rail infrastructure space and the rail service space should reveal how much robustness the service overlay supplies. Listed as hypotheses H1 to H4 in Table 4.4, these are expectations rather than conclusions. Their role is to set the prior, and the reader is invited to carry them into Chapter 5 and to read the empirical robustness curves there as confirmations, refinements or surprises.

These topological expectations enter the robustness analysis through the two edge orderings defined in Section 3.3.3. The first is the unweighted topology betweenness b^{topo}, computed on the structural supra-adjacency $\mathbf{A}^{\text{s}}$ whose topology this section has characterised, and the betweenness distribution of Figure 4.6 is its input. The second is the service-aware betweenness $\hat{b}$, computed on the supply layer E^{sup} by counting the origin-destination-departure buckets whose cost-minimising itinerary traverses each edge, and the enumerated pool of Section 4.3 is its input. The whole point of building the network as one formal object with a structural and a supply layer is that these two orderings act on the same edges, so the divergence between them in Chapter 5 is a result about the network rather than an artefact of two different datasets. The case study therefore supplies both inputs to that comparison, the structural topology here and the service-level pool in the next section.

Table 4.4: Topology-based disruption hypotheses H1 to H4, each with the layer it predicts is most affected and the topological mechanism behind it. The table is the bridge from this prior to the robustness analysis of Chapter 5, which tests each hypothesis.

	disruption family	most affected layer	mechanism
H1	single-edge failure	rail (L)	long diameter, low γ, no degree heterogeneity to absorb
H2	geographically correlated	rail (L) and intermodal at the same area	assortative corridor structure plus co-located cross-mode edges
H3	mode-specific shock	air alone, eased by the intermodal layer	scale-free air collapses, the layered network reroutes through rail
H4	targeted high-betweenness	air > intermodal > rail (L)	decreasing degree heterogeneity, heaviest at the air layer

4.2.4. Documented macro-hazards on the network

The two mechanism-based disruption families of Table 4.4, a mode-specific shock on the air layer and a geographically correlated shock on the rail layer, each have a well documented European instance, and the robustness analysis of Chapter 5 reads its macro-shock experiments against these two real events rather than against synthetic hazards.

The air-side instance is the April 2010 Eyjafjallajökull eruption. The eruption (14 April to 23 May 2010) lifted volcanic ash into the upper troposphere over a six-week period and triggered the largest peacetime closure of European airspace on record, with the cumulative operational impact reconstructed by Eurocontrol (2010) and Ulfarsson and Unger (2011) as approximately 100,000 cancelled flights across the peak-disruption phase. The eruption dynamics and the atmospheric transport that generated the cloud are documented by Sigmundsson et al. (2010), and the broader economic cost to European aviation and connected sectors is reconstructed by Mazzocchi et al. (2010). Because the closed airspace forms a contiguous geographic block rather than a uniformly random sample, it is the macro-scale instance of the mode-specific shock that the air layer of this network is least able to absorb.

The rail-side instance is the July 2021 Bernd storm. This extreme-rainfall episode flooded the Rhine, Ahr and Erft catchments of western Germany together with parts of Belgium and the Netherlands, with the meteorology characterised by Junghänel et al. (2021) and Mohr et al. (2023), the attribution of its rainfall intensity to anthropogenic climate change established by Kreienkamp et al. (2021) and Tradowsky et al. (2023), and the resulting damage to rail critical infrastructure documented by Koks et al. (2022). Its inundation area is mapped operationally by the Copernicus Emergency Management Service (Copernicus Emergency Management Service, 2021a, 2021b, 2021c), and it concentrates on the dense rail interior of the network, so it is the macro-scale instance of the geographically correlated shock that Table 4.4 predicts the rail infrastructure space is most exposed to.

4.3. Enumerated paths

The path-enumeration procedure of Section 3.2.5, applied to the supply layer E^{sup} of the functional multilayer network G^{f} over the three-year service window from January 2024 to December 2026, produces a unified intermodal path table of 482,947,150 alternatives organised into 12,581,130 origin-destination-departure buckets. The table is the empirical realisation of the supply layer, with each enumerated alternative a feasible sequence of scheduled supply edges, and it is the object on which the service-aware betweenness $\hat{b}$ of the previous section is computed. Applying the 24-hour trip-elapsed cap of Section 3.2.3 restricts this to the 429,952,440 alternatives across 11,068,358 buckets on which the robustness experiments are conducted, as Figure 4.8 shows. This section characterises what the enumeration produced, the redundancy it supplies to each bucket, the composition of the full table, and the interdependence of the cost-minimising itineraries. Figure 4.7 reports the first of these, the distribution of the choice-set size, namely the number of feasible alternatives available within a bucket. The median bucket carries 13 alternatives and the mean of 38.8 reflects a heavy tail that reaches 10,919, while six per cent of buckets carry a single alternative and therefore enter the disruption experiments with no enumerated redundancy.

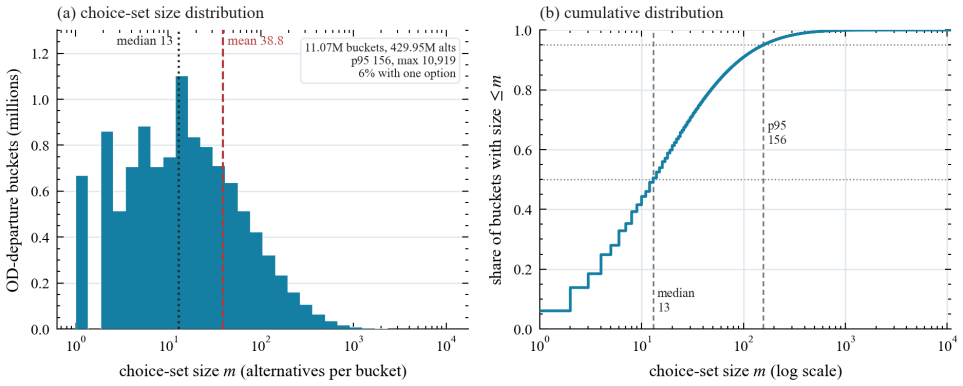

Figure 4.7: Distribution of the choice-set size, the number of feasible enumerated alternatives per origin-destination-departure bucket, across the 11,068,358 buckets carried into the robustness analysis. Panel (a) shows the distribution on a logarithmic size axis, with the median of 13 and the mean of 38.8 alternatives per bucket marked. Panel (b) shows the cumulative distribution, where half of the buckets carry at most 13 alternatives and 95 per cent at most 156, with a heavy tail reaching 10,919. Six per cent of buckets carry a single alternative and so offer no enumerated redundancy.

4.3.1. Mode-chain composition

The 482.9 million alternatives fall into thirteen distinct mode chains, reported in Table 4.5 and Figure 4.9. The thirteen are every sequence of rail and air legs of length one to three under the two-transfer cap, which is fourteen combinations, with the pure-air three-leg chain absent because the air pool is restricted to at most one flight leg per itinerary. Three-leg chains dominate, accounting for 95.07% of all alternatives. The single chain rail-air-rail accounts for 217,971,626 paths or 45.13% of all alternatives, and the all-rail chain rail-rail-

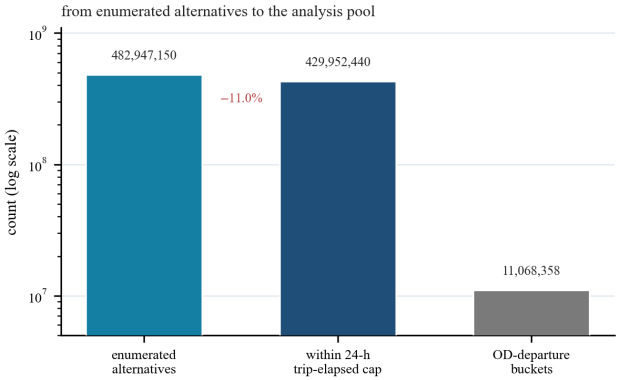

Figure 4.8: The full enumerated table of 482.9 million alternatives in 12.58 million buckets is reduced by the 24-hour trip-elapsed cap to the analysis pool of 429.95 million alternatives in 11.07 million buckets on a logarithmic scale.

rail for a further 110,722,004 paths or 22.93%. Pure-rail chains together sum to 123,201,824 paths or 25.51% of the table and pure-air chains to 357,471 paths or 0.07%. The remaining 359,387,855 paths or 74.41% are intermodal chains that use at least one rail leg and at least one air leg, and rail-air-rail alone is 60.6% of that intermodal population. Figure 4.10 summarises the same table by category and leg count.

The predominance of rail-air-rail has a structural origin in the coverage asymmetry of Section 4.1. The functional urban area framework of Section 3.1.1 places 377 origin and destination areas on the map, of which only 197 host an operating airport. For any pair whose endpoints are both areas without an airport, the only itinerary that uses a flight must begin and end on a rail leg with the flight in the middle, so rail-air-rail is the canonical air-using itinerary across most of the continent. The two-transfer cap then admits exactly one flight leg between two rail endpoints, and the per-bucket band filter admits the close-cost sister alternatives that differ only in feeder station, gateway airport or operator, which multiplies the count further. The same coverage asymmetry was already visible in Figure 4.4b, where 197 airports reach the whole continent that 1,674 stations serve on the ground.

4.3.2. Interdependence of the cost-minimising itineraries

The mode-chain composition counts every enumerated alternative. A stronger question asks, for each bucket, whether the cost-minimising itinerary uses both modes. Following the interdependence measure of Šfiligoj et al. (2026), let λ be the share of buckets whose minimum-cost alternative traverses both a rail and an air leg. On the 11,068,358-bucket analysis pool $\lambda = 0.543$. The per-bucket cost winner is rail-only in 44.2% of buckets, pure-air in 1.5%, and a true rail-air itinerary in 54.3%. Restricted to the 7,550,165 buckets for which an intermodal alternative exists at all, the intermodal share of winners rises to 0.796, and the remaining 3,518,193 buckets have no intermodal alternative and are served by rail alone. Figure 4.11 contrasts the raw composition with the winner composition and Figure 4.12 resolves the availability split behind the 0.796.

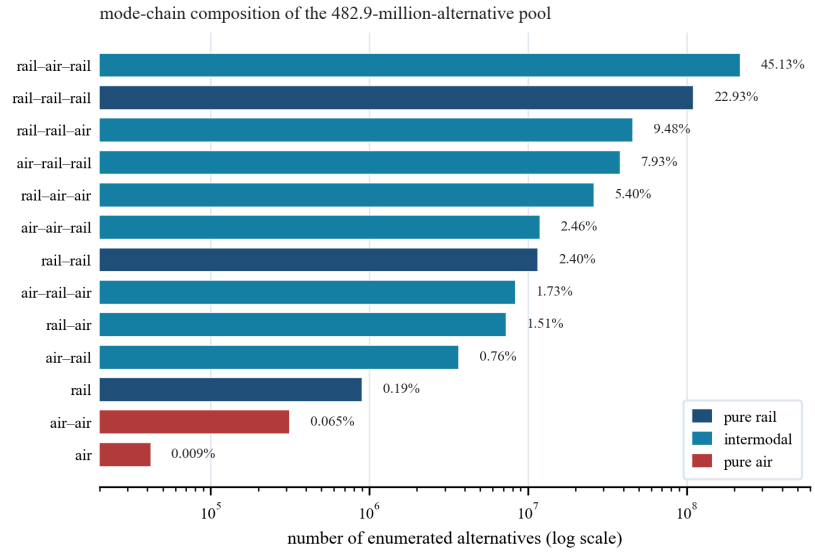

Figure 4.9: The thirteen mode chains on a logarithmic count axis, coloured by category. The pool spans more than four orders of magnitude, from rail-air-rail at 218 million to the single direct flight at 42 thousand, and the pure-air three-leg chain is absent by construction.

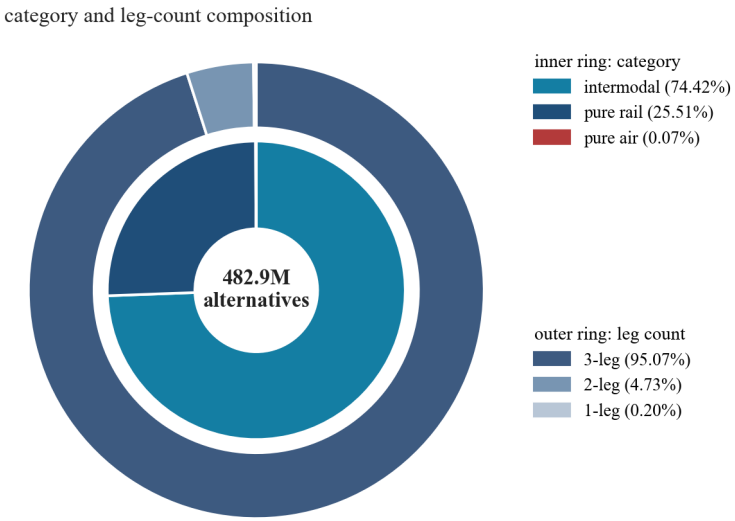

Figure 4.10: The same pool by category on the inner ring and by leg count on the outer ring. Intermodal chains are three quarters of the pool and three-leg chains are almost all of it.

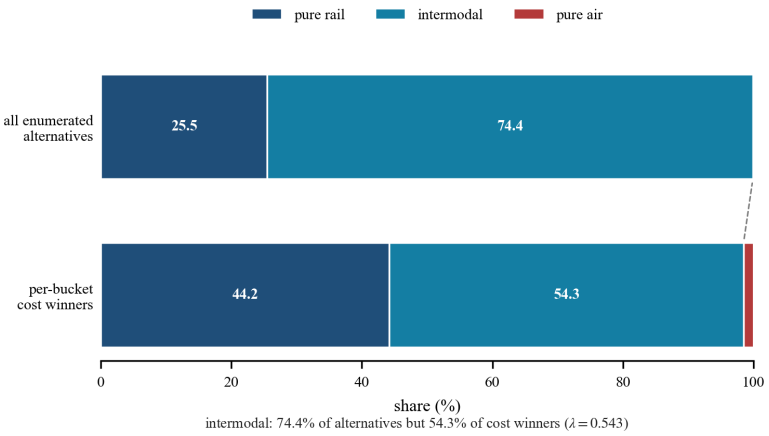

Figure 4.11: Category shares of all enumerated alternatives against category shares of the per-bucket cost winners. Intermodal chains are 74.4% of the alternatives but only 54.3% of the cost winners, because the more numerous intermodal alternatives are not always the cheapest in their bucket.

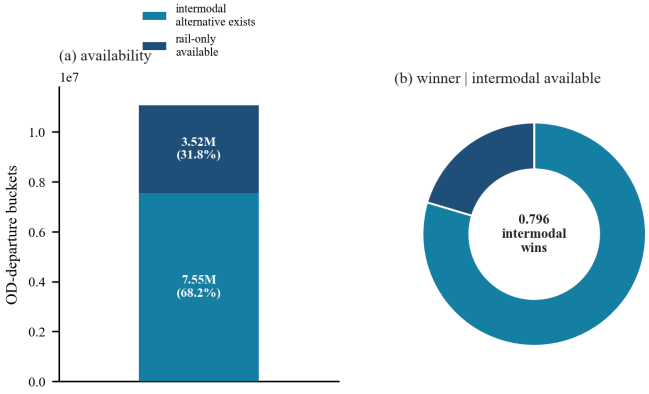

Figure 4.12: Panel (a) splits the analysis pool into the buckets that have an intermodal alternative and those served by rail alone. Panel (b) shows that among the buckets that have an intermodal alternative, the intermodal itinerary is the cost winner in 0.796 of cases.

Table 4.5: Mode-chain composition of the 482.9-million-alternative enumerated path table, grouped by category and ordered by share within each group.

Mode chain	Legs	Paths	Share
Intermodal (at least one rail and one air leg)			
rail-air-rail	3	217,971,626	45.13%
rail-rail-air	3	45,770,338	9.48%
air-rail-rail	3	38,313,473	7.93%
rail-air-air	3	26,102,817	5.40%
air-air-rail	3	11,903,788	2.46%
air-rail-air	3	8,372,723	1.73%
rail-air	2	7,289,874	1.51%
air-rail	2	3,663,216	0.76%
subtotal		*359,387,855*	*74.41%*
Pure rail			
rail-rail-rail	3	110,722,004	22.93%
rail-rail	2	11,579,487	2.40%
rail	1	900,333	0.19%
subtotal		*123,201,824*	*25.51%*
Pure air			
air-air	2	315,452	0.07%
air	1	42,019	0.01%
subtotal		*357,471*	*0.07%*
Total		**482,947,150**	**100.00%**

Two readings follow. First, the raw composition and the winner composition must be kept apart. Intermodal chains are 74.4% of all enumerated alternatives but only 54.3% of the per-bucket cost winners, because the abundance of intermodal alternatives does not make them the cheapest in their bucket. Second, an interdependence of 0.543 is high. The coupling between the rail and air layers is load-bearing in everyday cost-minimising routing rather than a disruption-only fallback, which supports treating the two networks as one interdependent system. The high value reflects the coverage asymmetry already noted, namely that the sparse air layer is reached from most origins only through a rail access leg, so air-using itineraries are predominantly intermodal.

4.4. Path-enumeration model validation

Before any robustness result is read, the path-enumeration model of Section 3.2.5 must be shown to reproduce published door-to-door travel times well enough to support network-level conclusions. The model precomputes, for every origin-destination pair and every feasible combination of legs with at most two transfers, the minimum generalised travel cost under the thesis cost formula, including intra-area connecting transfers. The benchmark is the rail-vs-air city-pair dataset of the European Commission (Brons et al., 2023; European Commission, Directorate-General for Regional and Urban Policy, 2022), which

reports total rail travel time for 1,356 directed European city pairs, matched to the thesis network by functional urban area name.

One scope difference must be stated at the outset. The benchmark reports home-to-home travel time, including access from the origin to the departure station and egress from the arrival station to the destination. The model reports station-to-station travel time only, with access and egress omitted by design as discussed in Section 3.2. A typical long-distance pair carries roughly twenty to fifty minutes of access and egress at each end, so the two datasets differ by a near-constant offset in favour of the station-to-station definition. A comparison of absolute minute values would therefore compare two different quantities. The validation is accordingly framed in terms of rank order and proportional agreement, which are invariant to a near-constant offset, rather than in terms of absolute error. For the same reason the thesis side is reported as the unweighted minute sum of in-vehicle time, platform waiting and intra-area connecting time, so that both datasets stay on a clock-minute scale.

Coverage is the first criterion. Of the 1,356 benchmark pairs, 1,266 could be matched to functional urban area pairs in the thesis network and yielded at least one enumerated path, which is the validation sample. The sample splits into 558 domestic and 708 cross-border pairs. On these 1,266 pairs the two datasets show strong rank-order agreement, with a Spearman correlation of 0.94, a Kendall correlation of 0.82 and a Pearson correlation of 0.84. These rank statistics are invariant to the access and egress offset and are the appropriate summary of agreement. Figure 4.13 shows the full diagnostics.

The ordinary-least-squares regression of the benchmark time on the model time has a slope of 0.66 and an intercept of 109 minutes. A slope below unity together with a large positive intercept is exactly what the home-to-home offset predicts, because the additive access and egress overhead is proportionally largest on short trips, where it lifts the benchmark time well above the model time, and smallest on long trips, where in-vehicle time dominates and the two times converge. The mean signed difference is a model time 24 minutes below the benchmark, consistent in sign and order of magnitude with the combined access, waiting and egress overhead that the benchmark applies at both ends. The mean absolute difference is 60.3 minutes and the median absolute percentage error is 18.7%, with 68.9% of pairs within one hour of the reference and 93.6% within two hours. Domestic and cross-border pairs agree comparably, with median absolute differences of 41 and 46 minutes respectively.

Exact minute-level validation against a matched station-to-station reference, for example official published in-vehicle schedules, is left to further research. The comparison reported here is sufficient to establish rank-order fidelity, which is the property the robustness analysis needs. A Spearman correlation of 0.94 means the set of pairs classified as fast or slow is preserved between the two datasets despite their definitional difference, so the robustness conclusions that depend on which stations, edges, corridors and operators are most critical rest on a sound empirical footing.

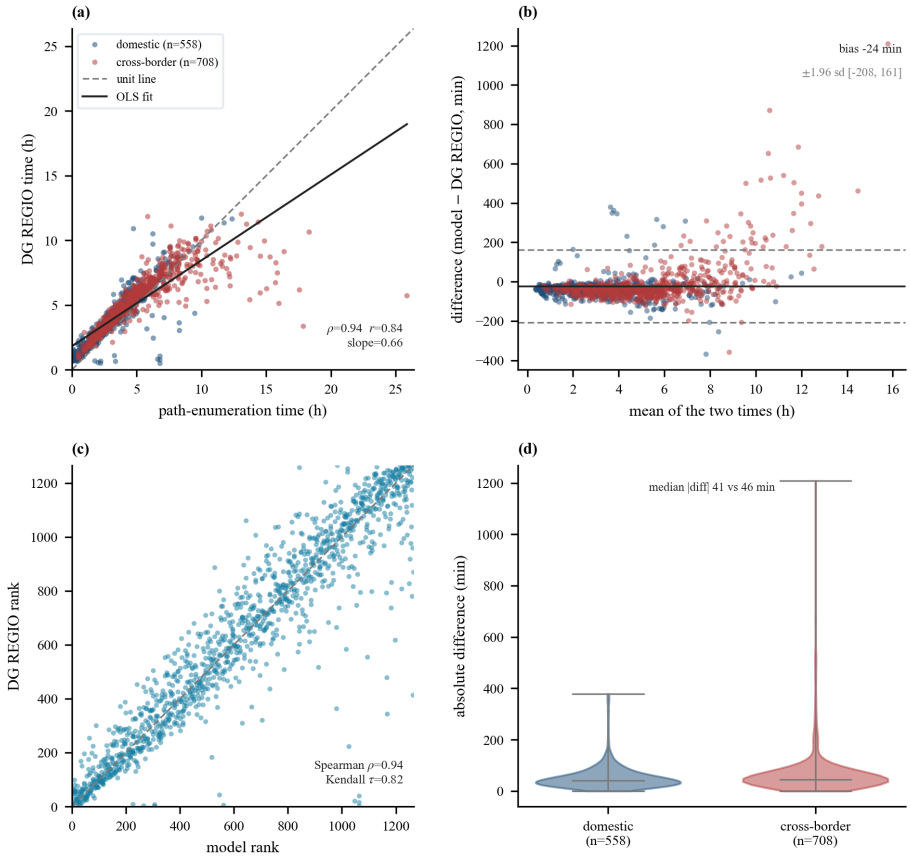

Figure 4.13: Validation of the path-enumeration travel-time model against the European Commission rail-vs-air city-pair dataset on 1,266 matched pairs. Panel (a) scatters the benchmark time against the model time with the unit line and the regression fit. Panel (b) is the agreement plot of differences against means. Panel (c) compares the two rankings. Panel (d) splits the absolute difference by domestic and cross-border scope.

5

Result

This chapter presents the results of the multilayer robustness assessment for the European long-distance rail and air service networks. Building on the passenger itineraries between European FUAs derived from the path enumeration algorithm of Section 4.3 and the disruption experiment framework of Chapter 3, the analysis evaluates how disruption scenarios affect the set of feasible passenger itineraries on the supply layer of the functional MLTN, rather than only the connectivity of the topological layer. These experiments act on the multilayer network constructed in Chapter 4.2 and on its 482,947,150 enumerated alternatives. Each experiment either gradually disables the network up to full removal of all edges, or disables part of it under a macro-shock scenario, and re-evaluates the surviving itineraries through the retained efficiency $R(x)$ and the associated robustness indices.

5.1. Type I: structural stress tests

Type I scenarios disable a fraction x of the supply-layer edges (consecutive-station rail segments and non-stop flights) under orderings based on different network centrality indicators or at random, and report the resulting passenger-level degradation on the supply layer of the functional MLTN. The removal order that incorporates service characteristics is the service-aware betweenness $\hat{b}(e)$ of equation 3.36, defined as the number of origin-destination-departure buckets whose alternative with minimal GTC value traverses edge e. Its agreement with two alternative orderings, the unweighted topology betweenness b^{topo} of equation 3.35 and a uniform-random baseline, is reported in Section 5.1.1. The extent to which the service-aware and topology orderings select the similar / dissimilar sets of edges is itself diagnostic of whether a topology-based criticality assessment can substitute for, or at best complement, a service-aware assessment on this network.

5.1.1. Betweenness-based failure

This section answers sub-question 2 of the thesis by examining how the quality of passenger itineraries degrades under the two edge-removal orderings. The major finding of this section is that the two rankings identify almost disjoint sets of supply edges on

the same layer of the functional MLTN, and the service-aware ordering produces higher retained-efficiency loss across different fractions of removed edges $x \in [0,1]$ than the topology ordering does. The cumulative robustness-loss integral $F = \int_0^1 [1 - R(x)]\,dx$ of equation 3.28 reaches $F = 0.934$ under the service-aware ordering against $F = 0.891$ under the unweighted topology ordering, a 4.8% excess concentrated at low edge-removal fractions. Additionally, the gap in the retained-efficiency ratio is already sharp at $x = 0.05$, where $R(x)$ falls to 0.168 under the service-aware ordering against $R(x) = 0.557$ under the topology ordering.

Betweenness-based ranking in the intermodal network

Prior to evaluating the network response to targeted failure, the difference between the two edge-removal orderings is first characterised. Topological and service-aware betweenness identify largely non-overlapping sets of critical edges in both the rail and air unimodal networks and in the intermodal network. This divergence is a direct consequence of whether or not the weight is assigned on basis of the possible itinerary provision situation from the perspective of passengers. The unweighted topology betweenness $b^{\text{topo}}(e)$ weights every station pair equally and counts all shortest paths, so it is largest on the structural bottleneck edges that carry the largest number of shortest paths between otherwise separated parts of the graph. The service-aware betweenness $\hat{b}(e)$ counts instead only the single minimal-GTC alternative of each origin-destination-departure bucket, so it is largest on the corridors that carry the heaviest itinerary load. Each ordering therefore elevates a different class of edge, the topology ordering the structural bottlenecks and the service-aware ordering the high-load corridors, and the sections that follow show where each class falls on the European map.

Rail edges dominate the upper tail of the intermodal edge rankings. In both the topology-based and service-aware rankings, the top-200 edges are exclusively rail edges. Thus, the most critical removable edges in the functional MLTN are located on the rail layer rather than on the air. However, the two rankings diverge not only in which edges they elevate but also in the spatial distribution of the edge set that governs robustness performance. The topology-based ranking mainly distributes importance across structurally pivotal bridge links wherever they arise in the network, whereas the service-aware ranking concentrates the critical edges within a smaller, geographically bounded core of heavily used operational corridors located predominantly within a single national network. Accounting for service characteristics therefore narrows the set of robustness-governing edges and localises it to a compact region, rather than dispersing it across the continent. The following two paragraphs examine each of these patterns in turn.

As seen in panel (a) of Figure 5.1 and Table 5.1, the unweighted topological ranking elevates edges whose removal fragments the structural MLTN. On the rail layer this generally produces edges that function as choke points between structurally distinct sub-networks, bridges that connect a weakly attached peripheral subnetwork to the continental core. The leading positions of the topology top-20 are dominated by edges that act as bridges between sub-networks, including the Channel Tunnel crossings (Brussel-Zuid to St. Pancras), the Franco-Spanish connections at the Pyrenees (Barcelona Sants to Perpignan and Lyon Part Dieu to Perpignan), the German-Polish connection (Berlin Hbf to Poznań Główny), the Alpine crossings (Chambéry to Milano, Zürich to Como), and the Belgian-German

and Belgian-Dutch connectors (Brussels to Cologne and Brussels to Amsterdam). Each of these edges lies on a large number of unweighted shortest paths because it provides one of the few direct rail connections across a sea, a mountain range, or an international border, but the later incorporation of service characteristics reveals that relatively few scheduled itineraries actually traverse through them. What is absent from the topological top-20 is the dense domestic rail corridor that the service-aware ranking elevates, and this absence reinforces the interpretation that unweighted betweenness rewards an edge's structural position rather than the number of scheduled itineraries that use it. The first air edge enters at position 1,154 on the Helsinki to Oulu flight, with a topology betweenness three orders of magnitude below the rail rank-1. Consideration of the air edges in the topological layer of the functional MLTN therefore does not redraw the structural cut-edge geography. Instead, it introduces a small set of peripheral air links, such as Erfurt–Palma de Mallorca, Brno–London Stansted, Helsinki–Oulu, and Madrid–Pamplona, which perform on the air layer a similar cut-edge function to that of the Pyrenees crossing on the rail layer.

In contrast, the service-aware ranking on the functional MLTN weights each edge by the number of scheduled itineraries that use it, and it selects a different set of critical edges, with no overlap between the two orderings of top-fifty rail edges. Panel (b) of Figure 5.1 shows the result. The service-aware ranking contracts edges with highest betweenness almost entirely inside Germany. Of the edges with top-20 highest betweenness, 19 entries lie on the German network and the Brussel-Zuid to Brussels North station-of-stations link is the sole non-German exception. The German concentration falls into four geographic clusters, the Rhine and Ruhr industrial corridor with Düsseldorf Hbf to Duisburg Hbf at rank-1 and Essen Hbf to Duisburg Hbf at rank-2, the Frankfurt-Mannheim spine led by Frankfurt (M) Hbf to Flughafen Fernbf at rank-3, the South German HSR corridors led by Ulm to Augsburg at rank-4 and Stuttgart to Ulm at rank-5, and a set of metropolitan station-of-stations through-links covering Hamburg, Munich, Göttingen to Kassel, and Berlin. Table 5.2 reports the full pattern partition with rank-by-rank assignments.

In sum, the two rankings differ as much in geography as in composition. Both place the most critical rail edges predominantly in western and central Europe, yet they disagree on where within that region those edges lie. The edges with high topological betweenness form a chain of international border crossings together with the edges from national networks that feed into them, appearing frequently along the frontiers stretching from the United Kingdom through Iberia, Switzerland, Italy, and Poland. The service-aware edges, by contrast, contract onto a compact set of short domestic corridors in Germany and its immediate neighbours. This contrast is visible as the tight cluster in panel (b) of Figure 5.1 against the edges which spread out more across continent in panel (a). The geographic concentration on the supply layer extends well beyond the strict top of each ranking and stays centred on Germany and its immediately neighbouring countries throughout the upper half. At top-200 the service-aware ranking is 88% German on the 4,108-edge rail layer and touches only four national networks (Germany, Belgium, the Netherlands, and France), against the topology ranking's 22.5% German and fourteen national networks. At top-500 the service-aware ranking holds 76% German with the remainder still on the Benelux-and-France core, while the topology ranking at 34% German already spreads across eleven additional national networks with 21% of its slots on cross-border edges. The two rank-

ings reach the 40.2% German baseline of the rail edge set only deep into their ranges, the topology at $K \approx 2{,}000$ and the service-aware at $K \approx 3{,}000$. This concentration on the German and Benelux core is therefore a genuine feature of the network and not limited to its top 20 edges. It reflects how the itineraries with minimum-GTC are routed through the German rail network and the short cross-border links around it. What is absent from the service-aware top-20 further confirms the divergence between the two rankings, since not a single international cut-edge appears among them. The Channel Tunnel, the Pyrenees gateway, and the Berlin to Poznań line all drop out, because none carries enough alternative-itinerary load to compete with the dense national corridors of the German core.

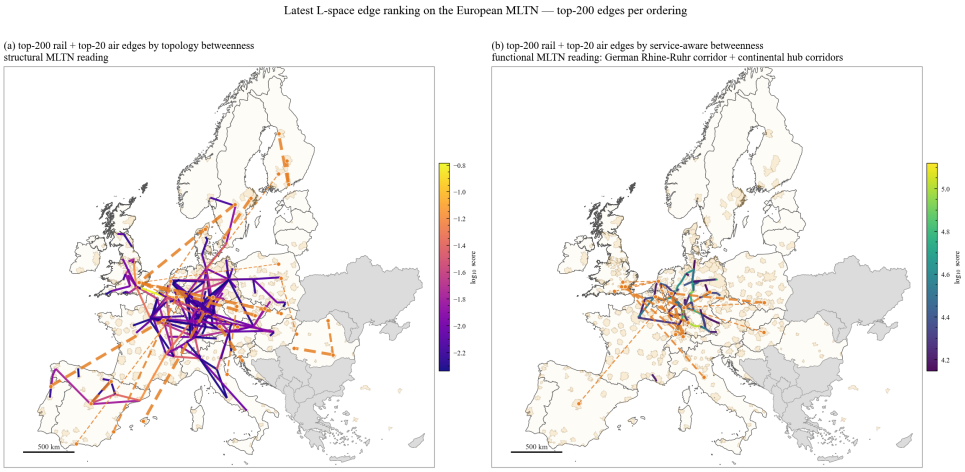

Figure 5.1: Edge ranking on the European MLTN, top-200 rail and top-20 air edges per panel. Panel (a) reports the topology betweenness ranking, with the leading edges concentrated on the international gateway shortcuts (Channel Tunnel, Brussels to Cologne, Iberian gateway, Berlin-Poznań, the Alpine crossings) and the peripheral air cut-edges. Panel (b) reports the service-aware betweenness ranking, with the leading edges concentrated on the German Rhine and Ruhr cluster, the Frankfurt-Mannheim spine, and the UK domestic feeders into London Heathrow. The two top-50 rail edge sets share zero edges, the two top-50 air sets share two. Excluded countries on both air and rail layer of the structural MLTN graph are hatched in red.

Edge ranking in the unimodal network

Since the intermodal edge ranking is dominated by rail edges at the top of both rankings, the rail unimodal ranking in the rail-layer of the structural MLTN is statistically near-identical to the intermodal one (Spearman $\rho = 0.93$ on topology and $\rho = 0.94$ on service layer) and therefore not reported separately for brevity.

The air-edge subset of the functional MLTN exposes the same topology-versus-service divergence that the rail-edge subset exposes, but with a partial overlap between the two betweenness rankings that the rail-edge subset does not show. The hub-and-spoke topology of the European air network across the 197 active airports already aligns the degree-one peripheral cut-edges and the high-itinerary-load trunk edges much more closely in graph terms than the rail-edge subset does, so the two air top-50 rankings share two common

Table 5.1: Topology-betweenness favours four patterns on rail. Rank in parentheses is the edge's position in the rail topology top-20.

Pattern	Concrete top-20 edges
Channel Tunnel and UK-London gateway chain	Brussel-Zuid–St. Pancras (1), St. Pancras–Finsbury Park (3), St. Pancras–Paris Nord (13), plus the St. Pancras–Selby (6), St. Pancras–Blackfriars (7), East Croydon–Blackfriars (11), Finsbury Park–Euston (10), and St. Pancras–Luton (19) spurs.
Iberian gateway through the Pyrenees	Barcelona Sants–Perpignan (4), Lyon Part Dieu–Perpignan (5), Barcelona Sants–Zaragoza (14).
Polish frontier and Eastern reach	Berlin Hbf–Poznań Główny (9).
Alpine and continental long-distance hops	Chambéry–Milano (16), Zürich–Como (20), Brussel-Zuid–München Ost (18), Chambéry–Brussel-Zuid (17), Lyon Part Dieu–Strasbourg (12), Paris Nord–Strasbourg (15), Cologne Hbf–Brussel-Zuid (2), Brussel-Zuid–Amsterdam Centraal (8).

Table 5.2: Service-aware betweenness favours four patterns on rail. Rank in parentheses is the edge's position in the rail service-aware top-20.

Pattern	Concrete top-20 edges
Rhine and Ruhr industrial corridor	Düsseldorf Hbf–Duisburg Hbf (1), Essen Hbf–Duisburg Hbf (2), Cologne Hbf–Düsseldorf Hbf (12), Dortmund Hbf–Bochum Hbf (13), Essen Hbf–Bochum Hbf (14), Cologne Hbf–Cologne Messe/Deutz (17), Cologne Messe/Deutz–Düsseldorf Hbf (18).
Frankfurt-Mannheim spine and Frankfurt airport access	Frankfurt (M) Hbf–Flughafen Fernbf (3), Flughafen Fernbf–Mannheim Hbf (10), Stuttgart Hbf–Mannheim Hbf (11), Frankfurt (M) Hbf–Mannheim Hbf (15).
South German HSR corridors	Ulm Hbf–Augsburg Hbf (4), Stuttgart Hbf–Ulm Hbf (5), München Pasing–München Hbf (7), München Pasing–Augsburg Hbf (16).
Metropolitan station-of-stations through-links	Hamburg Hbf–Hamburg Dammtor (6), Hamburg Hbf–Hamburg Harburg (8), Göttingen–Kassel Wilhelmshöhe (9), Brussel-Zuid–Brussels North (19, sole non-German entry), Berlin Hbf–Berlin Ostbahnhof (20).

edges (Aalborg–Amsterdam, Amsterdam–Münster/Osnabrück) against the rail-side zero overlap, and the two air top-20 rankings share one common edge (Aalborg–Amsterdam).

The service-aware betweenness on the unimodal air network elevates the feeder spokes that carry the most air itineraries into the main hubs, and the pattern reproduces across five geographic catchments. London Heathrow draws the UK domestic feeders (Aberdeen, Belfast City, Glasgow, Newcastle), whose only fast route into the long-haul network is air. London Stansted draws the continental and Irish low-cost feeders (Dublin, Cork, Hahn, Katowice, Rome Ciampino, Bergamo). Paris Charles-de-Gaulle draws continental feeders (Brest, Montpellier, Clermont-Ferrand) from regions the rail network does not serve by direct high-speed line. Amsterdam Schiphol draws Scandinavian and UK regional feeders (Billund, Aalborg, Birmingham) together with a spoke to Milan Linate. A fifth pattern, the Mediterranean island links (Madrid to Palma de Mallorca, Catania to Rome Fiumicino, Rome Fiumicino to Palermo), elevates routes where the sea removes the rail option entirely. As on the rail layer, the service-aware air ranking concentrates geographically rather than spreading evenly, here onto the two London hubs in particular, which appear in ten of the top-twenty edges.

The unweighted topology betweenness on the same edge subset elevates degree-one cut-edges that connect peripheral airports to the rest of the air network (Erfurt to Palma de Mallorca, Brno to Stansted, Helsinki to Oulu, Madrid to Pamplona). The top-11 edges share an identical betweenness score of 7.4×10^{-3} because each provides the only non-stop air link to its peripheral endpoint. These air cut-edges reach the top of the unweighted ranking for the same reason as the Channel Tunnel, the Pyrenees gateway, and the Berlin to Poznań line on the rail layer. Each is the single bridge linking a sparsely-connected part of the network to the rest, so a large number of the shortest paths between station pairs must pass through it, even though few scheduled itineraries actually do. These edges are therefore used by only a small share of the air itineraries despite their high topology rank, the same gap between shortest-path count and itinerary count that the two rail betweenness rankings reveal.

The reason the two air rankings disagree less than the two rail rankings is geometric. The hub-and-spoke topology of the European air network forces almost every shortest path between two non-hub airports to traverse at least one hub-to-hub link. The same hub-to-hub links also carry the largest share of fastest itineraries in the intermodal pool, so the two betweenness scores elevate overlapping sets of trunk edges on the air-edge subset, where the rail-edge subset's two scores pick up structurally distinct geographies instead. At top-200 the two air rankings overlap on 17.0% of edges (34 of 200) against the rail-side 4.0% (8 of 200); at top-500 the air overlap rises to 30.6% (153 edges) against rail's 14.2% (71 edges); at top-1000 the air overlap reaches 44.1% (441 edges) against rail's 24.0% (240 edges); and the full-pool Spearman rank correlation between the two betweenness rankings is $\rho = 0.18$ on the air-edge subset against $\rho = -0.03$ on the rail-edge subset. The two air rankings therefore agree weakly across the full edge set, whereas the two rail rankings are almost independent of each other. The divergence between an edge's structural importance and its itinerary load is present on both layers, but the hub-and-spoke geometry of the air network alleviates it throughout the ranking, not only in the top-20. The same geometry also drives the geographical concentration of the service-aware ranking onto a few hub airports, the air counterpart of the German concentration on the rail layer.

Table 5.3: Topology-betweenness on the air-only pool favours four patterns. Rank in parentheses is the edge's position in the air topology top-20. The top-11 edges share an identical betweenness score of 7.4×10^{-3} because each provides the only non-stop link between a peripheral airport and the connected component.

Pattern	Concrete top-20 edges
London Stansted peripheral feeders	Brno–Stansted (2), Örebro–Stansted (5), Stansted–Poprad-Tatry (7), Linz–Stansted (15), Ostrava–Stansted (19), Stansted–Olsztyn-Mazury (20).
National domestic feeders into the home hub	Helsinki–Oulu (3), Helsinki–Jyväskylä (6), Madrid–Pamplona (4), Ängelholm–Stockholm Arlanda (10), Bucharest Otopeni–Timişoara (11), Bucharest Otopeni–Suceava (12).
Holiday-route or charter cut-edges	Erfurt–Palma de Mallorca (1), Porto–Châlons-Vatry (8), Munich–Rijeka (9), Alicante–Maastricht (13), Prague–Rimini (14), Oradea–Warsaw (16), Málaga–Tampere (18).
Continental medium-traffic trunk (exception)	Aalborg–Amsterdam (17).

Table 5.4: Service-aware betweenness on the air-only pool favours five hub-anchored patterns. Rank in parentheses is the edge's position in the air service-aware top-20.

Pattern	Concrete top-20 edges
UK domestic feeders into London Heathrow	Aberdeen–LHR (1), Belfast City–LHR (2), Glasgow–LHR (3), LHR–Newcastle (4).
Low-cost spokes through London Stansted	Dublin–Stansted (7), Hahn–Stansted (12), Katowice–Stansted (16), Cork–Stansted (18), Rome Ciampino–Stansted (19), Bergamo–Stansted (20).
Continental feeders into Paris Charles-de-Gaulle	CDG–Montpellier (5), Brest–CDG (9), CDG–Clermont-Ferrand (15).
Amsterdam Schiphol spokes	Amsterdam–Billund (6), Aalborg–Amsterdam (8), Amsterdam–Birmingham (10), Amsterdam–Milan Linate (17).
Mediterranean intra-country trunks	Catania–Rome Fiumicino (13), Rome Fiumicino–Palermo (14), Madrid–Palma de Mallorca (11).

Network response during targeted failure

The topology and service-aware orderings of Section 5.1.1 identify almost disjoint sets of critical edges, and this section measures how the network responds when each is used to remove edges. Figure 5.2 reports three response indicators of Section 3.3.3, namely retained efficiency $R(x)$, mean mode-clipped reroute penalty $\Delta\bar{D}(x)$, and the disconnected share $f_{\text{disc}}(x)$ of OD-departure buckets, on the intermodal pool and on its pure-rail and pure-air constituents. Figure 5.3 pairs that view with a single removal-sequence map (left, edges coloured by the edge-removal fraction x at which each is removed) and the four-state outcome trajectory across the full edge-removal grid $x \in [0,1]$ (right). Together, the two figures address three questions about the network response, namely how far the quality of the feasible itineraries deteriorates, how the edge removals are distributed across the European network, and which of the four service states defined in equation 3.40 the affected OD-departure buckets fall into as the removal proceeds.

The retained efficiency ratio $R(x)$ measures how much of the network's itinerary quality each failure preserves. $R(x) = 1$ at $x = 0$ is the undisturbed network in which every OD-departure bucket retains its baseline minimum-GTC itinerary, and $R(x) \to 0$ is the highly fragmented and degraded network in which very few bucket retains any feasible itinerary. The early edge-removal range $x \in [0, 0.10]$ separates the two failure types most sharply. Under the service-aware betweenness, $R(x)$ falls from 1.000 to 0.168 at $x = 0.05$ and 0.100 at $x = 0.10$, while under the unweighted topology betweenness it falls only to 0.557 and 0.292 at the same two steps. Over this range the service-aware curve drops about twice as steeply, losing about 0.17 of $R(x)$ for each one per cent of edges removed, against 0.09 for the topology curve. The curve crosses $R(x) = 0.5$ at $x \approx 0.030$ under the service-aware betweenness and at $x \approx 0.061$ under the topology betweenness, so the service-aware failure halves the network's itinerary quality after removing about three per cent of edges, against six per cent for the topology failure.

Both failures converge to a common floor of $R(x) \approx 0.03$ by $x = 0.50$, where the service-aware curve reaches 0.032 and the topology curve 0.034 and about 98 per cent of OD buckets have lost every usable itinerary. The connections that remain reachable as edges are removed, however, differ markedly between the two orderings. At $x = 0.30$ the service-aware failure has removed almost all domestic German connections, leaving 1.3 per cent of German-internal pairs reachable against 3.4 per cent under topology. The self-contained

domestic networks of Italy and Spain, which the service-aware ordering reaches only late, retain about 4 per cent under the same failure, three times as much as German-internal pairs. The reason is that the service-aware ordering removes the highest-load rail corridors first, with the German trunk at the top of the order. Within the first tenth of edges ($x \in [0, 0.10]$) removed it also strips the domestic cores of the United Kingdom, Belgium, the Netherlands, and Switzerland, and a long international itinerary depends on several of these cores at once. Even connections with neither endpoint in Germany therefore keep a usable itinerary in only about 14 per cent of cases at $x = 0.10$, against 37 per cent under topology. The structural bridges removed first by the topology ordering carry comparatively few itineraries, with the top-20 carrying 2,593 minimal-GTC itineraries on average against 103,913 for the top-20 service-aware corridors, so the topology failure disconnects far fewer connections at every removal fraction short of the floor. In sum, two failures that reach the same $R(x)$ can still leave a very different set of connections in service, so the decision of which connections to protect cannot rest on $R(x)$ alone.

Additionally, the descent of the intermodal $R(x)$ curve toward this common minimum is constrained by the rail edges far more than the rail-only share of the baseline itineraries would suggest. Each OD-departure bucket is classed as rail-only or air-involving by the mode of its fastest baseline itinerary, and only 38.95 per cent of the 12,581,130 buckets are rail-only. Under the service-aware failure at $x = 0.05$, the rail-only buckets retain $R(x) = 0.095$ against 0.344 for the air-involving buckets, a factor of 3.6 in favour of the air-involving routing. However, 70.5 per cent of those air-involving buckets are already disconnected. The reason is that an air-involving itinerary cannot function without its rail-side airport-feeder edges. The rank-3 service-aware edge from Frankfurt Hbf to Flughafen Fernbf and the Frankfurt-Mannheim-Stuttgart spine are airport-access links that the rail-air-rail chain depends on, and removing them severs the air-involving itinerary at its rail-side feeder before the air leg can be reached. At low removal fractions, intermodal robustness is therefore set by the rail-side feeders rather than by the air legs that would otherwise supply it. Across the entire removal sweep the rail edges drive the intermodal $R(x)$ curve, and the air layer contributes only the robustness that survives the rail-side feeder constraint.

The right column of Figure 5.3 reports how the passenger status of each OD-departure bucket evolves as edges are removed. At every removal fraction x a bucket is unchanged from the baseline, reroutable within the GTC band, delayed beyond the band, or disconnected. Two operationally distinct response patterns emerge. Under the service-aware failure, the unchanged share collapses from 100 per cent to 16.0 per cent between $x = 0$ and $x = 0.05$, with the first 338 edge removals invalidating the baseline itinerary of 84 per cent of buckets. The delayed share stays negligible throughout the sweep (peak 0.87 per cent at $x = 0.05$). Affected buckets therefore pass almost directly from unchanged to disconnected, with no within-band detour in between. Under the topology failure, the unchanged share decays gradually (46.8 per cent at $x = 0.05$, 23.2 per cent at $x = 0.10$, 8.6 per cent at $x = 0.20$). The delayed share peaks at 3.4 per cent at $x = 0.10$ with a mean penalty of approximately 89 minutes, and the reroutable within-band share is also higher (9.7 per cent under topology against 4.6 per cent under service-aware at $x = 0.05$). The two failure types therefore impose qualitatively different operational problems. A topology failure leaves passengers delayed but still travelling, whom an operator can absorb through detour capacity and

delay-compensation procedures. A service-aware failure produces stranded passengers, for whom the only responses are replacement services or outright cancellation.

Although the random-failure benchmark is examined in detail in Section 5.1.2, one numerical observation from it is decisive here. Across the full removal range $x \in [0,1]$, the cumulative robustness-loss integral F under the unweighted topology betweenness sits just 0.36 standard deviations above the mean of the $M = 30$-seed random distribution of F ($z_F^{\mathrm{TOP}} = +0.36$), statistically indistinguishable from random allocation on the integrated criterion. The service-aware betweenness on the same range reaches $z_F^{\mathrm{SA}} = +9.20$, an extreme outlier against the same distribution. Figure 5.4 reports the z-trajectory of $R(x)$ at each removal fraction for both targeted failures against the random distribution. At the diagnostic removal fraction $x = 0.05$, the topology failure sits inside the random envelope at $z_R = -0.47$ ($p = 0.32$) while the service-aware failure sits well outside at $z_R = -8.51$ ($p < 10^{-16}$). The topology failure departs only mildly from random in the middle range ($|z_R| \approx 2$ at $x \in [0.10, 0.20]$), and both targeted failures cross back above the random envelope at $x \gtrsim 0.45$ as random removal continues to degrade $R(x)$. The integrated F statistic, the per-fraction z-trajectory, and the visual envelope comparison, all read from retained efficiency, therefore converge on the same operational consequence. In retained-efficiency terms, a protection budget allocated to the top of the unweighted topology betweenness on the intermodal edge subset is statistically indistinguishable from one allocated at random, a finding that the policy synthesis of Section 6.1 carries forward.

5

5.1.2. Random failure

This subsection reports the random edge-removal benchmark on the intermodal network. The disruption rule and the four indicators (retained efficiency $R(x)$, mean mode-clipped reroute penalty $\Delta\bar{D}(x)$, disconnected share $f_{\mathrm{disc}}(x)$, and the cumulative robustness-loss integral F) are defined in Section 3.3 of the methodology chapter. Each configuration is averaged over $M = 30$ Monte-Carlo seeds, a replicate count justified in Section 5.4.3 and Figure 5.15.

The degree-distribution test of the case study found neither layer to be cleanly scale-free, with the rail layer statistically exponential and the air layer hub-dominated but cut off at high degree (Table 4.3). The strong tolerance to random failure that a scale-free structure classically confers (Albert et al., 2000) is therefore not assured here. A five percent uniform-random edge loss disconnects 37.6 per cent of OD-departure buckets and brings retained efficiency to $R(x) = 0.577$, a loss of 42.3 per cent with no part of the network preferentially targeted. Averaged across the 30 seeds, random failure yields a robustness-loss integral of $F = 0.889$ (inter-seed standard deviation 0.005), almost exactly the topology failure's $F = 0.891$ and below the service-aware failure's $F = 0.934$. Random failure degrades the network almost as heavily as the topology-betweenness ranking, despite choosing its edges with no preference.

Random failure is also the structural baseline against which the targeted-failure curves of Section 5.1.1 are read. The random-failure was introduced into complex-networks analysis by Albert et al. (2000) as the error counterpart to directed failure, and chapter 8 of Barabási (2016) reviews the foundational asymmetry that scale-free networks are robust to such errors but vulnerable to targeted failures. Mattsson and Jenelius (2015) translate this distinction to transport-robustness research as the non-deliberate disruption category

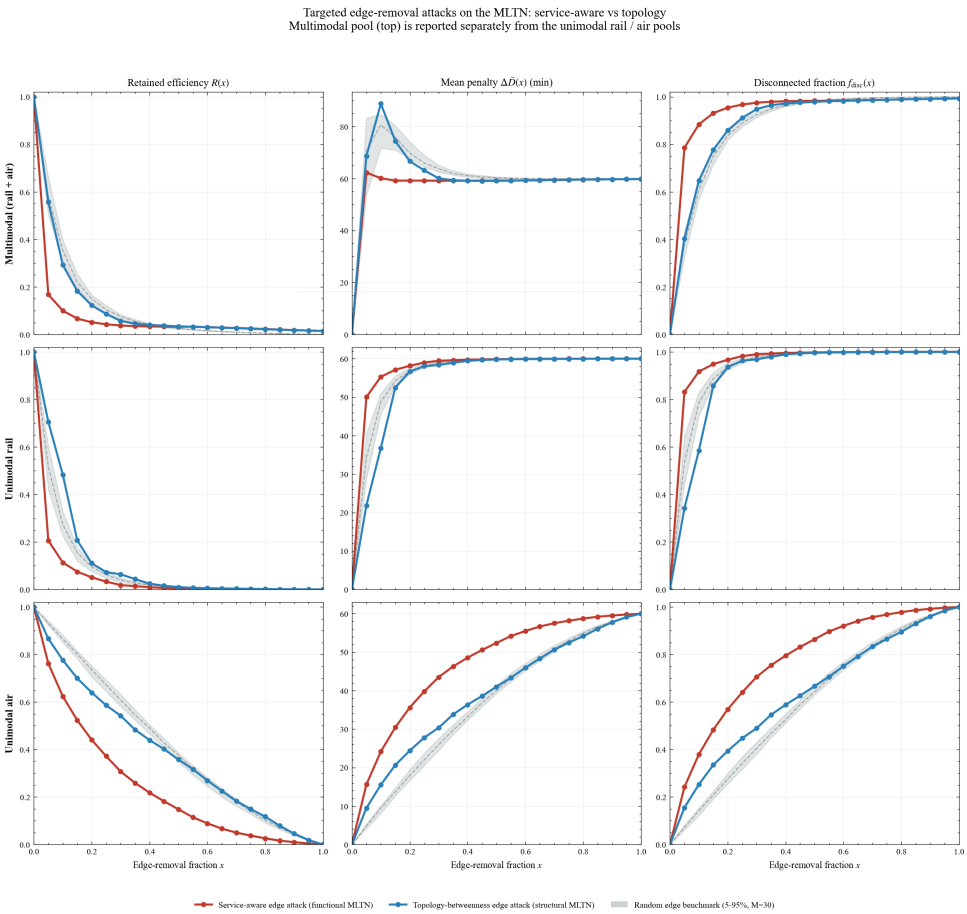

Figure 5.2: Targeted edge-removal failures on the intermodal pool (top row) and its unimodal rail and unimodal air sub-pools (middle and bottom rows). Columns report retained efficiency $R(x)$, mean mode-clipped reroute penalty $\Delta\bar{D}(x)$, and disconnected bucket share $f_{\mathrm{disc}}(x)$ across the edge-removal fraction $x \in [0, 1]$. The service-aware ordering (red) drops the intermodal $R(x)$ to 0.168 at $x = 0.05$, against 0.557 under the topology ordering (blue) and 0.577 under the random benchmark (grey band, 5th to 95th percentile across $M = 30$ seeds). The two targeted curves diverge sharply at low edge removal fraction and converge to a common floor at $R(x) \approx 0.03$ by $x \approx 0.30$.

that operators absorb daily; on the European network it has been studied for the rail layer by Bešinović (2020) and for the air-traffic layer by Cardillo et al. (2013), with scattered weather events that strike many countries simultaneously as the dominant operational realisation. The major European long-distance rail operators and airlines face this baseline as the level of network degradation for which their routine planning is sized.

Anatomy of random failure

Random failure is conceptually unlike the two targeted-failure rankings of Section 5.1.1. The two betweenness-based rankings each encode a fixed ordering on the edges of the

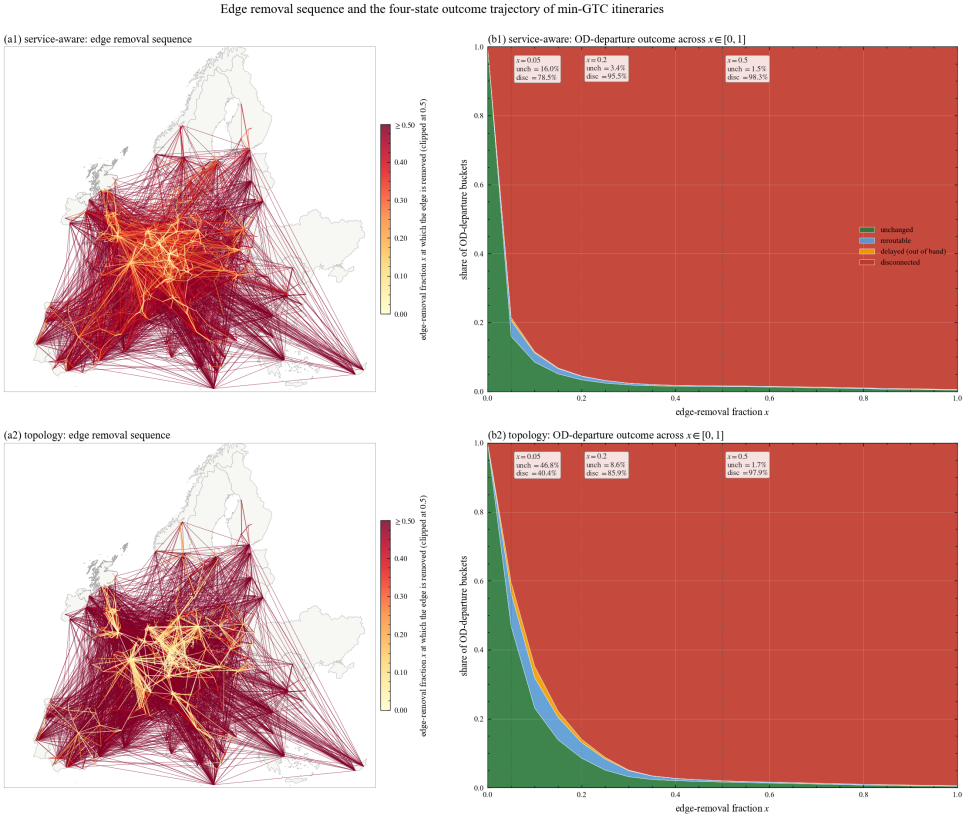

Figure 5.3: Left column reports the removal sequence on the intermodal rail layer as one map per ordering. Each edge is coloured by the edge-removal fraction x at which it is removed, from pale cream for the earliest removals through orange to deep red for edges removed at $x \geq 0.5$. Under the service-aware ordering the earliest removals, in cream and orange, form a compact cluster centred on Frankfurt that covers the Rhine and Ruhr corridors, the Frankfurt-Mannheim-Stuttgart main line, and the airport-feeder edges the rail-air-rail chain depends on, while the long peripheral connections in deep red are removed last. Under the topology ordering the earliest removals are instead dispersed across continental cut-edges, the Channel Tunnel, the Pyrenees gateway, the Alpine crossings, and the German-Polish connection. The right column reports the passenger status of OD-departure buckets across the full removal range $x \in [0,1]$, the share of buckets in each of the four mutually exclusive states (unchanged baseline, reroutable within the GTC band, delayed beyond the band but still structurally connected, fully disconnected). Annotations at $x \in \{0.05, 0.20, 0.50\}$ give the unchanged-baseline and the fully-disconnected share at each step.

functional MLTN, with the topology betweenness elevating structural cut-edges and the service-aware betweenness elevating corridors that carry many fastest-GTC alternatives. The random benchmark involves no ordering at all. Each of the $M = 30$ Monte-Carlo seeds draws a single uniform permutation of the edge set, and the removal-fraction grid $x \in \{0.05, 0.10, \ldots, 1.00\}$ removes nested prefixes of that permutation. The resulting outcome distribution is reported as an inter-seed envelope rather than as a single degradation trajectory from one failure scenario.

The map view of Figure 5.5 shows how the random benchmark differs from the two

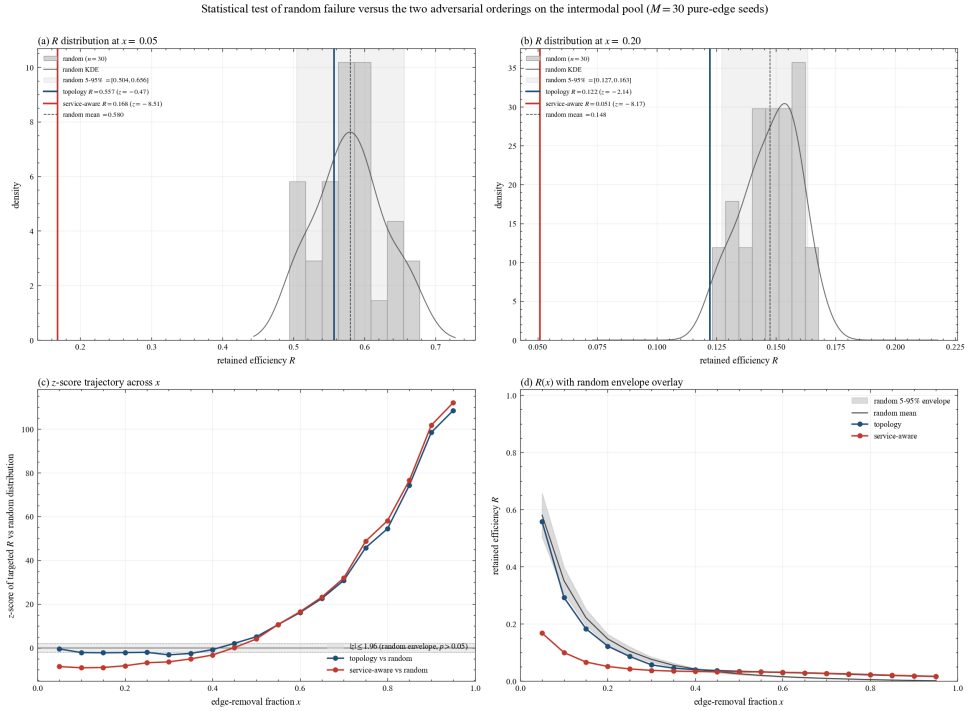

Figure 5.4: Statistical test of random failure against the two betweenness-targeted failures on the intermodal pool ($M = 30$ pure-edge seeds). (a) Distribution of random retained efficiency $R(x)$ at $x = 0.05$, with topology and service-aware values overlaid (topology at $z_R = -0.47$ inside the random envelope, service-aware at $z_R = -8.51$ outside). (b) Same at $x = 0.20$ where topology has moved just outside the lower fifth percentile. (c) Full z_R-trajectory of both targeted failures across $x \in [0, 1]$; shaded band marks $|z| \leq 1.96$, namely the random envelope at the 5% two-sided significance level. (d) Retained-efficiency curves with the random 5th-95th-percentile envelope overlaid.

targeted failures. The left column reports, for every edge, the median (top) and the mean (bottom) edge-removal fraction x at which it is dropped across the 30 seeds. Both maps are spatially uniform, with every edge's median and mean removal fraction clustered near the midpoint $x \approx 0.5$ and no edge removed systematically earlier or later than any other. This even spread is the visual signature of uncorrelated disruption, in contrast to the concentration of the targeted maps of Figure 5.3.

The network's performance degrades substantially even under this uniform failure, because the itinerary load it acts on is concentrated disproportionately on a small share of the corridors. Across the rail corridors, the top five per cent carry 34.6 per cent of the itinerary load and the top tenth carry 51.5 per cent, while the most traversed single corridor lies on the path of 150,279 OD-departure buckets against a median of 2,092. These corridors carry the most itineraries, and those itineraries have few alternatives. A uniform draw samples this concentration rather than avoiding it, removing the heaviest corridors at the same rate as any other edge, so that at five per cent removal 37.6 per cent of OD-departure buckets lose every itinerary while only 9.7 per cent reroute within the band.

Random failure therefore differs from the service-aware ordering in timing rather than reach. That ordering removes the heaviest corridors first while a uniform draw removes them at scattered points, but either way they are removed, so random failure still strips the operationally heaviest parts of the network rather than sparing them. Those parts are the western German corridors that the service-aware ranking of Section 5.1.1 placed at the top of the network's critical edges, set apart by both their compact geography and their disproportionate, poorly-substituted itinerary load. The structural stress tests therefore locate the network's robustness on this one region whether a failure targets it or strikes at random.

Operationally, the random benchmark represents the class of hazards whose affected area is uncorrelated with the chapter's two criticality rankings. Storm Ciara (February 2020) and Storm Eunice (February 2022) are recent illustrations of weather systems that struck many countries simultaneously without concentrating on a single corridor (Met Éireann, 2020; Mühr et al., 2022), the dominant operational realisation of random failure in the rail-layer review of Bešinović (2020) and the air-layer modelling of Cardillo et al. (2013). The benchmark is a reference distribution, not an upper bound on real-world damage, because spatially correlated events such as the July 2021 Bernd floods (Section 5.3.2) can concentrate their removals on the same geography the service-aware betweenness elevates and therefore exceed the random envelope at the same edge count.

Network response during random failure

Figure 5.5 pairs the removal-sequence maps with the four passenger status under the performance degradation curve across the full edge removal percentage $x \in [0,1]$. The right column reports the trajectory aggregated across the 30 seeds as a single central line, namely the median (top row) and the mean (bottom row), with the seed-to-seed 5th to 95th percentile envelope on the unchanged share shown as a dashed black line. Three findings emerge.

The retained-efficiency ratio curve $R(x)$ under random failure falls from $R(x) = 1.000$ at $x = 0$ to median $R(x) = 0.577$ at $x = 0.05$ (5th to 95th inter-seed envelope $[0.504, 0.656]$, mean $\Delta\bar{D} = 69.7$ minutes), to median $R(x) = 0.354$ at $x = 0.10$, to $R(x) = 0.220$ at $x = 0.15$, to $R(x) = 0.150$ at $x = 0.20$, and to $R(x) = 0.025$ at $x = 0.50$. The initial slope is approximately 8.5 percentage points of retained efficiency per percentage point of edge-removal budget, and the half-life of $R(x)$, the edge-removal fraction at which $R(x)$ first crosses 0.5, falls at $x \approx 0.067$. The median and the mean track each other within one to two percentage points across the entire grid, which confirms the Monte-Carlo convergence of the 30-seed benchmark.

The four types of passenger status under the network degration curve of random failure are also dominated by disconnection rather than by delay. At $x = 0.05$ the median seed reports 48.9% of OD-departure buckets unchanged, 9.7% reroutable within the GTC band, 3.4% delayed beyond the band but still structurally connected, and 37.6% fully disconnected. By $x = 0.10$ the unchanged share falls to 27.0% and the disconnected share rises to 60.6%, while the reroutable and delayed bands together never exceed 14% of the OD-departure set across the entire removal range. The same operational signature seen under the two targeted-failure rankings holds under random failure as well. The European intermodal network does not provide a smooth degradation pathway in which most

disrupted itineraries find a slower substitute, and the dominant mode of operational decay is outright disconnection of the OD pair regardless of how informed the failure protocol is. This is a structural property of the network's cluster geometry rather than of any particular ordering.

The comparison to the two targeted-failure curves of Figure 5.2 produces a key operational finding of this subsection. At $x = 0.05$ the topology betweenness leaves $R(x) = 0.557$ against the random median $R(x) = 0.577$, a difference of only 0.020 retained-efficiency units that sits inside the random envelope $[0.504, 0.656]$ at z-score -0.47, statistically indistinguishable from a uniform random sub-sample of the same number of edges. The service-aware betweenness leaves $R(x) = 0.168$ at the same removal fraction, a 0.409-unit gap below the random benchmark at z-score -8.51, roughly twenty times the topology-versus-random gap. From $x \approx 0.10$ onward the topology betweenness drops just below the lower fifth percentile of the random envelope and the two curves separate slowly, yet the topology-versus-random gap stays small at every removal step relative to the service-aware-versus-random gap. Unweighted graph-theoretic betweenness is therefore not informative for a failure protocol on this network, and the classical random-versus-targeted distinction collapses unless the targeting ordering is itinerary-aware.

5

5.1.3. Policy insights from the structural stress tests

Targeted failure

The service-aware failure in Section 5.1.1 concentrates the network's critical edges on Germany's national rail network, particularly the Rhine-Ruhr corridor around Düsseldorf, Duisburg, Essen and Cologne, the Frankfurt–Mannheim spine, and the high-speed lines toward Stuttgart and Munich. Witlox et al. (2022) describe the Ruhr area as an intensively used network on which high-speed services lose operating speed, and Otsuka et al. (2017) characterise the Rhine-Alpine axis through Düsseldorf, Cologne, Frankfurt and Mannheim as a capacity-constrained mixed passenger and freight corridor. Three further lines of evidence place the criticality in the same region. Five of the nine TEN-T core network corridors cross Germany, more than any other member state (German Federal Ministry of Transport, 2026). The Rhine-Alpine axis, the most heavily used mixed rail corridor on the continent, carries its longest bottlenecks in its German section, with the Frankfurt-Mannheim spine and the Karlsruhe-Basel approach as the principal constraints (Nollert & Niedermaier, 2019). Deutsche Bahn attributes its long-distance punctuality of 63.4 per cent in the first half of 2025 to overload at the Frankfurt, Cologne and Hamburg hubs, whose delays transfer to the rest of the network (Deutsche Bahn AG, 2025). Computed from timetable and itinerary data alone, the service-aware ranking converges on the same corridors that the TEN-T policy map, the corridor congestion record and the operator's own diagnostics independently identify. European transport policy treats these corridors as vitally important and invests heavily in their capacity, with the German Deutschlandtakt programme committing decades of targeted infrastructure funding to them (Cosandey & Ropelius, 2026).

However, current European transport policy rarely treats the traffic concentrated on these corridors of the German national network as a robustness problem in its own right. Because that network carries a disproportionate share of Europe's cross-border itineraries, as demonstrated in Subsection 5.1.1, a failure on these corridors is not confined to Ger-

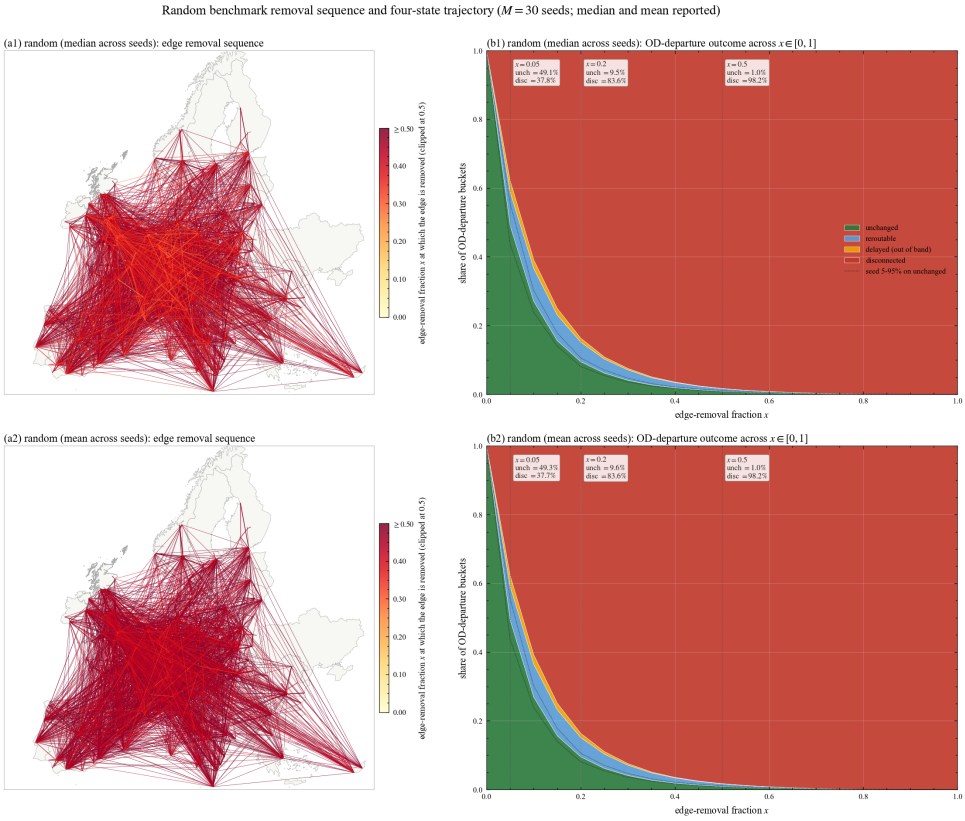

Figure 5.5: Left column reports the per-edge median (top) and mean (bottom) edge-removal fraction x across the $M = 30$ pure-edge random seeds on the same edge pool as the importance-based failures. The near-uniform deep-red colour across the entire European map is the visual signature of uniform sampling, namely that no edge is systematically removed earlier or later than any other across seed replicates. Right column reports the four-state Cats trajectory aggregated across the 30 seeds as a single central line (median in the top row, mean in the bottom row), with the seed-to-seed 5th to 95th percentile envelope on the unchanged share shown as a dashed black line. Three reference verticals at $x \in \{0.05, 0.20, 0.50\}$ annotate the unchanged and fully-disconnected shares at each intensity step. The median and mean trajectories track within one to two percentage points across the grid, which confirms the Monte-Carlo convergence of the 30-seed benchmark.

many, and most of the additional damage the service-aware failure inflicts beyond a topology failure falls on cross-border journeys, which lose their fastest feasible itinerary outright rather than finding a slower substitute. The recent record shows this spillover directly. In 2017 the collapse of a tunnel under construction at Rastatt closed the Karlsruhe–Basel line for seven weeks and severed the Rhine–Alpine corridor, with measured delay propagating onto the Swiss network hundreds of kilometres from the site, at an estimated loss above two billion euros (Büchel et al., 2020; HTC Hanseatic Transport Consultancy and European Rail Freight Association (ERFA) and Network of European Railways (NEE) and UIRR, 2018). In 2025 Swiss Federal Railways began terminating two Deutsche Bahn EuroCity services at Basel, severing two direct cross-border connections rather than con-

tinue to import the disruption from Germany (The Local, 2025). The German national network therefore sets the robustness of the wider system, and a neighbouring operator has already cut a through-service rather than continue to absorb the disruption that crosses the border.

Two consequences follow for the allocation of a robustness budget. The same budget protects far more retained efficiency on the German trunk corridors than if it were spread evenly or placed on the international crossings, because, measured in retained efficiency, placing it by topological betweenness is statistically indistinguishable from placing it at random at the same intensity ($z_R = -0.47$ against the random distribution of Figure 5.4). No air service absorbs a failure on these corridors either, since carriers do not fly the short distances these corridors span (Bleijenberg, 2020), so the air substitute that a purely topological reading would count as available is rarely usable in practice (Figure 5.6). The concentration also makes German domestic policy a continental matter, because Central Europe reaches the western network largely through Germany and the international traffic of the Benelux, France and Switzerland runs to and through it. German domestic pricing and maintenance decisions therefore fall on those countries' cross-border services as much as on its own, as the 2026 Court of Justice ruling against the German charging formula and the retreat of new open-access services to domestic routes both show (International Railway Journal, 2026; RailTech, 2025). The case for treating the German network as a European rather than a national responsibility rests on its robustness, not only on its capacity.

The divergence between service-aware and purely topological perspectives shown in Section 5.1.1 indicates that a robustness assessment should not lean too heavily on topology alone. A purely structural ranking stays the appropriate tool where service information is unavailable, or where one abstraction must compare networks of different scale and geography, and complex-network analysis is valued for exactly that generalisability across disciplines (Barabási, 2016; von Ferber et al., 2009). Where service information is available, however, incorporating it can shift the picture a topological analysis provides. On the European MLTN the structural ranking favours the cross-cluster cut-edges of Section 5.1.1, which carry less itinerary traffic and keep more alternative paths around them, as shown by the cross-modal substitutes of Figure 5.6, and it overlooks the within-cluster German rail corridors that carry the most feasible itineraries and have no usable substitute. A protection budget set from pure topology alone is therefore directed away from the edges that matter most once service characteristics are layered onto the same graph, the misallocation that Bešinović (2020) and the Frankfurt-over-Paris strength reranking of Calzada-Infante et al. (2020) both identify once frequency and itinerary weight are restored.

Policy insight of random failure

A scale-free network withstands random failure far better than a deliberate one, its connectivity resting on hubs too rare for a blind draw to strike, so that only a targeted removal reaches them (Albert et al., 2000; Barabási, 2016). That robustness is conditional rather than generic, requiring a heavy-tailed power-law degree distribution, the structure that concentrates connectivity on the few hubs a random draw is unlikely to reach, and it disappears for the exponential degree distributions of ordinary networks such as an electrical

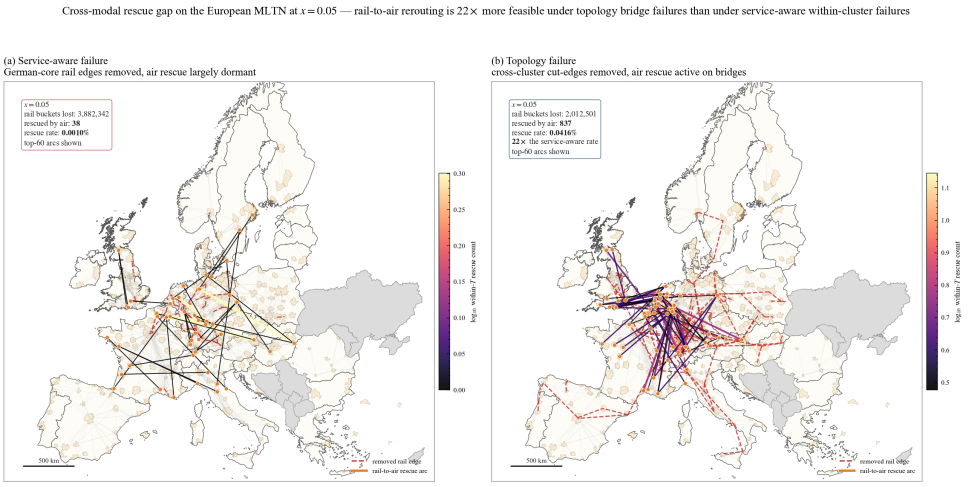

Figure 5.6: Cross-modal rescue gap on the European MLTN at $x = 0.05$. Each panel draws the top-200 removed rail edges as dashed red lines and, as solid arcs, the rail-to-air reroutes available to the disrupted itineraries, with arc colour the $\log_{10}$ within-tolerance rescue count and orange dots the endpoint airports. Countries with neither active rail nor air service are shaded grey. (a) Under the service-aware failure the removed edges concentrate on the German Rhine-Ruhr and Frankfurt-Mannheim trunks and the cross-modal substitute is largely dormant, with 3,882,342 rail-baseline buckets lost and only 38 rescued within tolerance, a rescue rate of 0.0010 per cent. (b) Under the topology failure the removed edges fan out across the cross-cluster cut-edges such as the Channel Tunnel, the Iberian gateway and the Alpine crossings, and the substitute is active on the bridges, with 2,012,501 buckets lost and 837 rescued within tolerance, a rescue rate of 0.0416 per cent and 22× the service-aware rate. The top 60 arcs are shown per panel.

power grid (Barabási, 2016). The European intermodal network is that same exponential case in the layer that dominates its supply layer, the case study having found its rail layers exponential in degree with no hub for a random draw to spare (Chapter 4.2). Its air layer is the one genuinely scale-free layer, but that tail is cut off at high degree and air is the smaller of the two modes, so it cannot lend the combined network the random tolerance the scale-free result describes. Section 5.1.2 then shows that on the integrated criterion, the cumulative robustness-loss integral F, random failure is statistically indistinguishable from the unweighted-topology targeting ($z_F = +0.36$) and is exceeded only by the itinerary-aware failure ($z_F = +9.20$), because the concentration of itinerary load means a blind draw still strikes the few heavily used German edges the service-aware ranking elevates. The accidental disruptions it represents, for instance, the winter storms that cross several countries at once and the signalling and rolling-stock failures operators meet as routine (Bešinović, 2020; Mattsson & Jenelius, 2015; Met Éireann, 2020; Mühr et al., 2022), are already familiar from daily operation, tracked through punctuality and reliability reporting and handled by the disruption-management procedures built for them. What the analysis adds is the network-level total those incident-by-incident records cannot show, that an uncorrelated draw of such failures does damage on a par with a deliberately targeted one, which argues for treating them in network-wide contingency planning rather than as isolated local incidents.

Because that failure falls evenly across the map (Figure 5.5) and cannot be located in

advance, the response it calls for is not to protect particular edges, the lever the targeted failure pointed to, but to hold the reserve capacity that Cats and Jenelius (2015) identify as the value of planning for the unexpected. They show that replacement services and reserve vehicles and crew are the real-time response to an unforeseen disruption, and that without such spare capacity the demand rerouted onto the surviving network exceeds it and cascades into denied boarding and wider welfare loss. That reserve should be sized to the upper edge of the seed envelope, the 95th-percentile outcome at $R(x) = 0.504$ with close to half of OD-departure buckets disconnected, because an uncorrelated draw can fall on those same German edges and approach the targeted damage, and it should be held for re-accommodating stranded travellers at the affected endpoints rather than for detours, since the damage falls as outright disconnection rather than recoverable delay. Beyond these operator lessons, the benchmark serves the reference role the vulnerability literature assigns to random failure (Albert et al., 2000; Ippolito & Cats, 2024; Pien et al., 2015), the null of no prioritisation against which every targeted, operator, and macro-shock scenario in this chapter reads as a degree of exposure rather than a bare loss of efficiency.

5.2. Type II: supply-layer withdrawal under operator removal

Type II scenarios are applied at the operator level, removing the supply-layer edges an operator provides in the functional MLTN rather than individual topological-layer edges. The methodology of Section 3.1.4 attaches an operator attribute $\pi(e)$ to every supply-edge $e \in E^{\text{sup}}$ through the per-service attribute tuple of equation 3.13, and the operator-withdrawal protocol of Section 3.4.2 removes every supply-edge satisfying $\pi(e) \in \Pi^{\text{withdraw}}$ for a designated operator subset $\Pi^{\text{withdraw}} \subseteq \Pi$. The set of operators analysed here is the union of 51 rail train operating companies and 31 air carriers, 82 operators in total. Each operator's full set of supply-layer edges is removed as one unit, which makes the single-operator withdrawal a single-point-of-failure test in the reliability-engineering sense. The indicator suite is then recomputed on the remaining intermodal network to measure how far the network's service depends on a single operator.

Two variants are reported. Operator withdrawal applies the edge-removal protocol of Section 3.3.3 at operator granularity and is reported in three cases, $k = 1$ single-operator, $k = 2$ pairwise among the ten most damaging single-operator withdrawals, and $k = 3$ triplet withdrawal among the largest operators. Cooperation-tier subgraph restriction, the second variant in Section 3.4.2, is instead defined by which operators a passenger may use, retaining an itinerary only when every leg is run by a member of a defined cooperation set C.

This section deepens the tension established in Section 5.1 by identifying the operator behind the service-aware concentration. The headline empirical finding is that the Deutsche Bahn $k = 1$ withdrawal alone leaves the intermodal network at $R(x) = 0.390$, the loss a random or topology-ranked failure produces at about 8 per cent of edges removed and a service-aware failure at about 4 per cent. The B1-to-B5 cooperation gradient of 0.624 retained-efficiency units then exceeds the service-aware-versus-random infrastructural gap of 0.409 units by a factor of 1.53×. The mechanism is that the service-aware German interior of Section 5.1 is also the set of supply-layer edges Deutsche Bahn serves

and the binding constraint on the cooperation-tier subgraph, so the edge-level and the operator-level assessments identify the same German corridors as critical. This reinforces the chapter's central finding that topology-based and service-based assessments diverge, and shows that here the divergence is carried by one identifiable operator rather than by network geometry alone.

5.2.1. Single-operator withdrawal

The $k = 1$ enumeration removes each of the 82 operators in the European intercity service network (51 rail train operating companies and 31 air carriers) in turn and records the resulting bucket-level $R(x)$ on the intermodal pool. Each withdrawal removes the supply-layer edges $\{e : \pi(e) = \pi^\star\}$ for the withdrawn operator $\pi^\star$ on the supply layer of the functional MLTN of Section 3.1.4, using the per-edge operator attribute of equation 3.13, while the underlying topological-layer L-space edges remain intact. Supply-layer edges for which the withdrawing operator is the sole provider lose all feasible service and no longer support any itinerary; supply-layer edges that other operators continue to serve survive at reduced service intensity.

The full catalogue of the 82 single-operator withdrawals is reported in Table 6.3, sorted by retained efficiency. Deutsche Bahn is the single most damaging operator withdrawal on the intermodal networks and its provision of feasible itineraries, at $R(x) = 0.390$ and a disconnection share of 54.4%; more than half of the OD-departure buckets lose their fastest feasible itinerary when DB's supply-layer edges are removed. A second tier of national rail operators (NS, SNCF, NMBS/SNCB, DFTO, ÖBB) leaves the network at $R(x) \in [0.890, 0.949]$ when withdrawn, disconnecting between 4.4% and 12.9% of OD-departure buckets. Ryanair and easyJet are the only air carriers in the top ten most damaging operator withdrawals, at $R(x) = 0.950$ and $R(x) = 0.971$ respectively. Below the top ten the catalogue tail decays rapidly, with 61 of the 82 operators at $R(x) \geq 0.99$ and 45 at $R(x) \geq 0.997$, so most single-operator withdrawals leave the network's retained efficiency essentially intact. Two operators sit at exactly $R(x) = 1.000$ because they serve no edge of the intermodal network once the study's geographical coverage is applied. Hellenic Train has no scheduled service on Greek rail network and Icelandair's only catchment lies outside the FUA coverage, while the remaining 80 operators each carry at least one supply edge. The asymmetry at the head of the distribution is not a property of operator size alone. SNCF serves 407 supply-layer edges against DB's 1,499, a quarter as many, yet its withdrawal leaves $R(x) = 0.919$ against DB's 0.390, roughly an eighth of the damage. The damage an operator inflicts reflects the itinerary load its edges carry.

Figure 5.7 positions each operator on the same edge-removal axis $x \in [0, 0.30]$ that the random benchmark and the two betweenness-based rankings of Section 5.1.1 traverse. The edge-equivalent of an operator withdrawal is the count of supply-layer edges the operator serves. Deutsche Bahn serves 1,499 supply-layer edges, corresponding to $x = 0.222$ on the edge-removal axis, the largest single-operator edge count in the catalogue and comparable in size to Ryanair (FR) with 1,364 edges at $x = 0.202$. No other operator approaches that share, with easyJet (U2) at $x = 0.071$, SNCF at 0.060, and ÖBB at 0.056, while the remaining top-ten operators are all below $x = 0.045$. The figure labels only the headline scenarios inline, with the full 82-operator distribution visible as the faint background scatter and tabulated in appendix Table 6.3. The per-edge concentration of the operator-withdrawal

damage is markedly lower than what either betweenness-based ranking would inflict at the same edge count. The DB withdrawal at $R(x) = 0.390$ sits substantially above the service-aware curve at $x = 0.22$ ($R(x) \approx 0.05$), above the topology curve ($R(x) \approx 0.09$), and above the random median ($R(x) \approx 0.14$). DB's supply-layer edges therefore span both the heavily-traversed German main lines that the service-aware ranking of Section 5.1.1 elevates (the Rhine-Ruhr corridor, the Frankfurt-Mannheim spine, the southern German high-speed lines, and the Frankfurt central-station-to-airport feeder) and a far larger set of regional and secondary edges of low itinerary load, so that a small high-load subset accounts for most of the damage.

The DB withdrawal forms the second line of evidence in this chapter that the edges representing rail corridors with the greatest impact on the European intermodal network's robustness are concentrated within Germany's national rail network, where a small number of corridors carry a disproportionate share of the feasible itineraries. The service-aware edge ranking of Section 5.1.1 places the Rhine-Ruhr corridor, the Frankfurt-Mannheim spine, and the South German HSR alignments in the top twenty rail edges of the intermodal pool. Both lines of evidence converge on the rail networks within Germany's national borders, the chapter's strongest evidence so far that the network's robustness, measured by the quality and travel time of passenger itineraries, depends on a small set of service corridor rather than on the cut-edge types that pure topological betweenness identifies across the graph as a whole. The macro-shock scenarios of Section 5.3 test whether a regional weather hazard over those same corridors offers further confirmation. A robustness-enhancement programme with a geographical focus on these corridors would therefore preserve the passenger service that depends on them, since they carry the largest share of the network's feasible itineraries. It would not, however, preserve the graph's topological connectivity, which rests on a separate and almost disjoint set of international cut-edges, so one geographically focused budget cannot protect both service and connectivity at once. And because rail-air co-existence supplies little usable substitution at the short distances these corridors span, the service they carry has no second-mode fallback if they are left unprotected.

Table 5.5: Ten most damaging single-operator withdrawals on the intermodal pool, sorted by retained efficiency $R(x)$ (most damaging first). The full 82-operator catalogue is reported in appendix Table 6.3.

Operator	Mode	Edges	x	$R(x)$	f_{disc} (%)	$\Delta\bar{D}$ (min)
DB	rail	1499	0.222	0.390	54.4	60.8
NS	rail	200	0.030	0.890	12.9	18.8
SNCF	rail	407	0.060	0.919	11.5	14.8
NMBS/SNCB	rail	143	0.021	0.919	8.0	10.3
DFTO	rail	291	0.043	0.936	6.6	9.3
ÖBB	rail	377	0.056	0.949	4.4	4.7
FR	air	1364	0.202	0.950	9.7	19.1
SBB	rail	269	0.040	0.956	7.5	12.3
Trenitalia	rail	260	0.038	0.961	5.2	7.5
U2	air	479	0.071	0.971	5.4	11.7

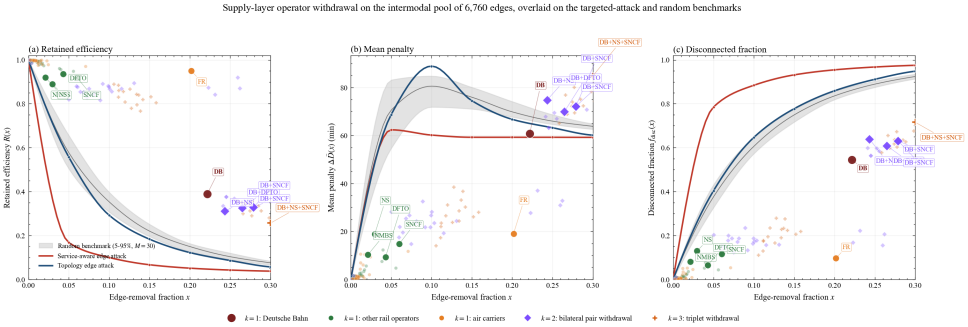

Figure 5.7: Supply-layer operator withdrawal on the intermodal pool of 6,760 edges, overlaid on the targeted-failure and random benchmarks. Panel (a) reports retained efficiency $R(x)$, panel (b) reports the mean mode-clipped reroute penalty $\Delta\bar{D}(x)$, and panel (c) reports the disconnected bucket share $f_{\text{disc}}(x)$. Background curves are the service-aware failure in red, the topology failure in blue, and the random benchmark as the grey envelope (5-95th percentile across $M = 30$ seeds) with the random median in dark grey. Operator-withdrawal scenarios are overlaid as scatter markers, red filled circles for Deutsche Bahn at $k = 1$, green for the other $k = 1$ rail withdrawals, amber for the $k = 1$ air carriers, purple diamonds for the $k = 2$ bilateral pairs, and orange plus signs for the $k = 3$ triplets. Each marker is positioned at the edge-removal fraction x. Headline scenarios are labelled inline; the full catalogue is in appendix Table 6.3.

5

5.2.2. Pair and triplet withdrawal

Since network robustness under failure is more sensitive to DB than to any other operator, the pairwise and triplet withdrawals quantify the additional damage that further operators contribute. The $k = 2$ case evaluates every pair among the ten most damaging single operators and reports the worst of these together with seven named scenarios such as the joint withdrawal of DB and ÖBB. Figure 5.7 overlays the resulting retained-efficiency, mean-penalty, and disconnected-fraction values on the same edge-removal axis as the $k = 1$ withdrawals. The most damaging pair is DB and NS at $R(x) = 0.311$, combining the German domestic network and the Dutch long-distance network in a way no other pair reproduces, followed by DB and DFTO at $R(x) = 0.326$ and DB and SNCF at $R(x) = 0.328$. Among triplets the most damaging is DB, NS, and SNCF at $R(x) = 0.256$, with the ten worst triplets falling within $R(x) \in [0.26, 0.33]$. The damage each added operator contributes diminishes rather than compounds. The second and third operators together lower $R(x)$ by only about 0.13 beyond DB alone, from 0.390 to 0.256, and of that reduction the second operator accounts for 0.079 and the third for only 0.055. This diminishing pattern reflects the concentration of service-aware betweenness among a few operators, DB foremost, so that once the most damaging operators are removed each additional one contributes progressively less.

The disconnected share in panel (c) of Figure 5.7 rises with each added operator, from $f_{\text{disc}} = 0.544$ under DB alone to 0.639 under the worst $k = 2$ pair (DB plus NS) and 0.716 under the worst $k = 3$ triplet (DB plus NS plus SNCF), increments of 9.5 and 7.7 percentage points whose diminishing sequence matches that of the retained-efficiency losses. The two panels record the same transition of buckets between outcome states. At each step the share of unchanged baselines declines by nearly the same amount as the disconnected share increases, while the share of buckets that retain a within-tolerance reroute

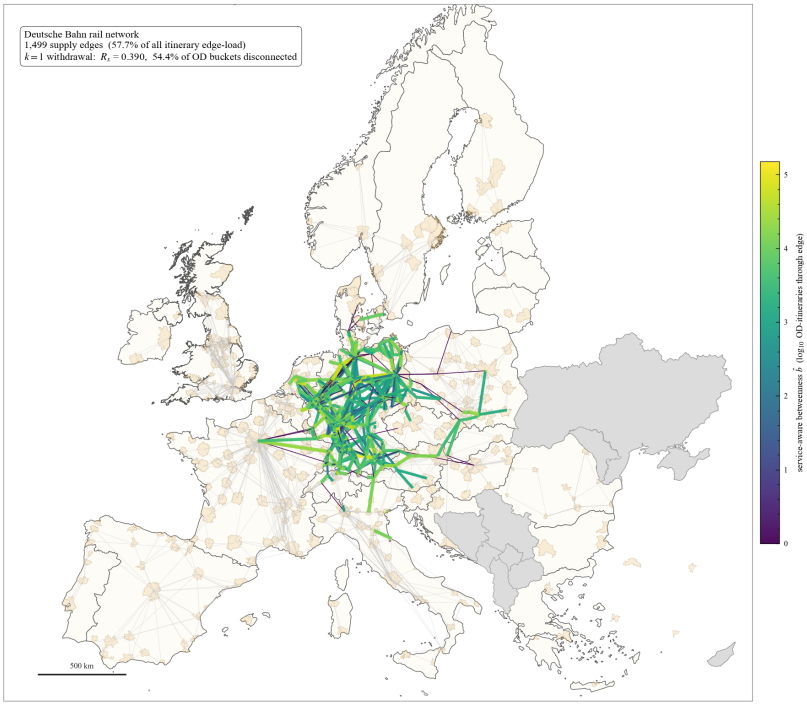

Figure 5.8: The Deutsche Bahn rail network on the European MLTN, with its 1,499 supply edges coloured and weighted by service-aware betweenness $\hat{b}$, the number of OD-departure buckets whose minimum-GTC itinerary traverses each edge. The remaining rail edges are drawn in faint grey for context. Deutsche Bahn reaches broadly across Germany and carries 57.7 per cent of the total passenger itinerary, and that load concentrates on the west-central German corridors of the Rhine-Ruhr area and the Frankfurt-Mannheim spine rather than spreading evenly across the edges it serves. This concentration is why the single-operator withdrawal of Section 5.2.1 is the most damaging in the catalogue, removing the cheapest itinerary of 54.4 per cent of OD-departure buckets and leaving retained efficiency at $R(x) = 0.390$.

remains between 1 and 2 per cent, so each added operator converts previously unaffected OD-departure buckets directly into disconnections rather than lengthening the delays of the buckets that remain served. The mean-penalty panel (b) distinguishes the operator withdrawals from the edge-removal benchmarks. The mean reroute time $\Delta\bar{D}$ rises from 60.8 minutes at DB $k = 1$ to 74.7 minutes at the worst $k = 2$ pair and 84.8 minutes at the worst $k = 3$ triplet, while the service-aware, topology, and random benchmarks remain between 59.2 and 67.3 minutes at the corresponding intensities, because at these removal fractions the benchmark orderings disconnect more than 85 per cent of buckets entirely and their mean penalty is therefore computed over a small residual set of rerouted itineraries. The DB withdrawal remains within this benchmark band, whereas the pair and triplet withdrawals exceed it, producing longer detours per affected bucket than any of the three benchmarks at the same removal fraction. The rerouted itineraries follow more circuitous paths because the operators that substitute for the withdrawn services are peripheral to the German main lines.

5.2.3. Cooperation-tier subgraph restriction

The cooperation-tier experiment inverts the withdrawal operation. Instead of disabling a set of operators, the surviving alternative set is restricted to itineraries whose every leg's operator lies in a specified cooperation set C. Five named tiers are evaluated, and they correspond to identifiable cooperation arrangements in current European long-distance practice. B1-strict is the single-operator subgraph that captures the passenger holding a single-operator ticket. B1-split allows exactly one operator change. B2 is the seven commercially documented bilateral pairs such as DB and ÖBB on the Nightjet, DB and SNCF on the Paris-Frankfurt-Stuttgart corridor, DB and SBB on the Munich-Zürich corridor, DB and NMBS on the Aachen-Liège axis, SNCF and Renfe on the Madrid-Lyon axis, SNCF and Trenitalia on the Paris-Milan corridor, and Eurostar after the post-Thalys restructuring. B3 is the Railteam alliance of nine cooperating rail operators. B4 is the full rail-interline subgraph of 51 rail TOCs, the MITA analogue at continental scale. B5 is the full modal-interline subgraph of all 82 rail and air operators, which corresponds to the no-disruption baseline.

Figure 5.9 reports the $R(x)$ progression across the tiers, the cooperation-set size $|C|$ underneath each bar, and the four-state Cats outcome distribution per tier. The B1-strict tier restricted to Deutsche Bahn alone retains $R(x) = 0.376$ on the intermodal pool, which is operationally near-symmetric with the Deutsche Bahn $k = 1$ withdrawal of Section 5.2.1. DB withdrawal removes DB's supply-layer edges and keeps the other 81 rail and air operators available for booking, leaving $R(x) = 0.390$ functional. The B1-strict-best setting keeps DB's supply-layer edges available for passengers and removes the supply layers of other 81 operators, leaving network with $R(x) = 0.376$. The two settings report essentially the same retained efficiency from opposite directions, the operator-side and the booking-side of the same supply-layer geometry. The cooperation-tier curve then rises from $R(x) = 0.376$ at B1-strict-best through $R(x) = 0.426$ at B1-split (allowing one operator change to the best second carrier, ÖBB on the Nightjet) and the B2 bilateral pairs at $R(x) \in [0.39, 0.43]$ to $R(x) = 0.631$ at the B3 Railteam alliance, $R(x) = 0.733$ at B4 full rail interline, and $R(x) = 1.000$ at B5 full modal interline. The B1-to-B5 institutional gradient of approximately 0.624 retained-efficiency units exceeds the service-aware-versus-random infrastructural gap of approximately 0.409 retained-efficiency units at the $x = 0.05$ edge removal fraction, so how far the passenger is allowed to rebook matters more than the infrastructural budget of the worst itinerary-aware failure process, and is the binding constraint on European long-distance robustness.

The narrow bilateral pairs in the B2 set deliver markedly less than the DB-anchored ones. SNCF and Renfe on Madrid-Lyon reach only $R(x) = 0.022$, SNCF and Trenitalia on Paris-Milan reach $R(x) = 0.028$, and Eurostar after the post-Thalys restructuring reaches $R(x) = 0.0005$. These narrow pairs cover too small a share of the OD-departure pool to deliver meaningful robustness, since their cooperation scope is bounded geographically to a single international city pair. Only the DB-anchored bilateral pairs in the B2 set reach $R(x) \approx 0.4$, because Deutsche Bahn's domestic German network already covers the high-betweenness rail corridors on which most through-itineraries route. Bilateral cooperation agreements alone therefore do not substitute for alliance-tier interlining and only the Railteam-tier set or broader delivers operationally measurable retained efficiency.

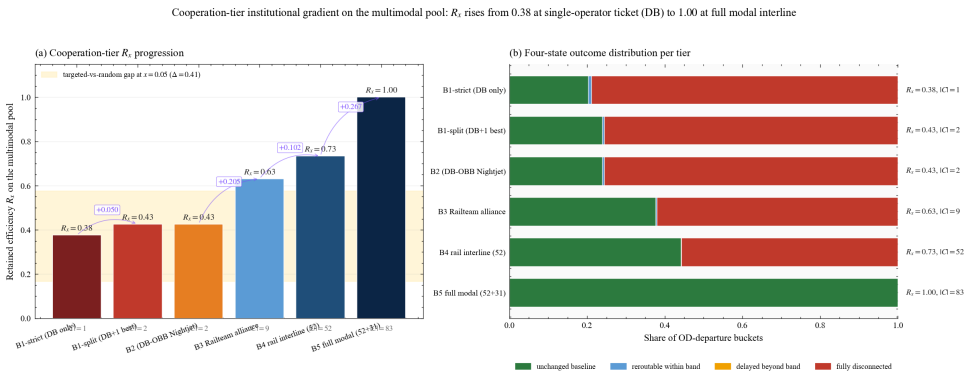

Figure 5.9: Cooperation-tier institutional gradient on the intermodal pool. Panel (a) reports the retained efficiency $R(x)$ across the six named tiers (B1-strict single-operator ticket, B1-split one-change, B2 bilateral pair, B3 Railteam alliance, B4 full rail interline, B5 full modal interline). The cooperation-set size $|C|$ is annotated below each bar, and the step-by-step dividend over the previous tier is annotated above. The shaded yellow band marks the service-aware-versus-random infrastructural gap of approximately 0.41 retained-efficiency units at $x = 0.05$, namely the largest infrastructural gap measured in Section 5.1.1. The B1-to-B5 institutional gradient of 0.624 units exceeds this band, namely how far the passenger is allowed to rebook is the binding constraint on European long-distance robustness. Panel (b) reports the distribution of four passenger status for the same tiers as horizontal stacks, showing how cancelled trips convert progressively to unchanged baselines as more operators join the cooperation set.

5.2.4. POLICY INSIGHT

Where the stress tests of Section 5.1 examined the network on the infrastructure side by removing individual edges from its structure, the supply-layer withdrawal and cooperation-tier experiments contribute the supply side of the chapter's answer to sub-question 3 on the extent to which rail-air co-existence enhances robustness on the European long-distance network. Both are defined on the supply layer of the functional MLTN of Section 3.1.4, where every service edge carries the operator that provides it. Those stress tests showed that an air alternative that is feasible purely in terms of topological connectivity is rarely usable in practice, because carriers do not fly the short distances the critical rail corridors span, and the macro-scenarios of Section 5.3 find the reverse substitution, from air to rail, also absent under a continent-scale airspace closure. The cooperation-tier gradient provides the booking-side counterpart. Restricting itineraries to a single operator's services suppresses retained efficiency as severely as the worst single-operator withdrawal, and extending them across the 82 rail and air operators restores it stepwise to the undisrupted baseline. Rail-air co-existence can therefore enhance robustness only where passengers may rebook across operators and modes. The implications that follow are presented for three parties, the EU regulator, the operators, and the research community, each addressing one component of sub-question 4 on the most critical service links and corridors and the targeted measures that respond to them.

For the EU regulator, the cooperation-tier curve identifies cross-operator and cross-modal rebooking as the dominant institutional levers on the network. The Railteam-tier transition (B1-to-B3) closes 0.255 retained-efficiency units and the cross-modal transition (B4-to-B5) adds a further 0.267, together accounting for 84% of the cooperation-tier divi-

dend. European Parliament and Council of the European Union (2021) provides limited cross-operator rights today and treats rail and air as separate regulatory systems and Toet et al. (2026) document the practitioner consensus that this integration remains unrealised, and Bruno et al. (2025) quantify the air-rail substitution potential at continental scale. The empirical record argues for strengthening EU 2021/782 into a binding cross-operator re-booking default under disruption and for treating rail and air as jointly bookable substitutes, in line with the cross-border passenger-rail action plan of European Commission (2021). The second implication concerns where protection is needed. Two diagnostics that act on the different layer of function MLTN, the service-aware ranking of Section 5.1.1 and the Deutsche Bahn $k = 1$ withdrawal of Section 5.2.1, identify the same German corridor, from the Rhine-Ruhr area to Frankfurt-Mannheim, as the part of the network on which the most feasible itineraries depend. The corridor's cities appear among the routes on which rail is already faster than air in the door-to-door mapping of Brons et al. (2023), and Witlox et al. (2022) name the Ruhr among the intensively used networks where operating speeds constrain high-speed rail. The next TEN-T programme cycle (European Parliament and Council of the European Union, 2024b) should therefore add an instrument dedicated to protecting this concentration of service, alongside the cross-border connectivity the programme already prioritises.

From the perspective of the operators, two operational implications follow. First, Deutsche Bahn is the operator on which the largest share of the network's service depends. Every one of the most damaging pair and triplet withdrawals of Section 5.2.2 involves Deutsche Bahn together with a neighbouring national rail operator whose international itineraries run to and through the German network. Removing Deutsche Bahn from the network and restricting passengers to it alone leave almost the same retained efficiency (Section 5.2.3), so the dependence is structural rather than commercial. O'Sullivan and Patel (2004) describe the underlying problem, that when transport operations are fragmented among many separate companies, the integrity of the system as a whole becomes the responsibility of none of them. On the European network measured here, however, that integrity depends in effect on a small set of operators, Deutsche Bahn foremost. It is therefore essential that these operators adopt a whole-system perspective in their operational and commercial decisions. The experiments, moreover, understate Deutsche Bahn's prominence, since they capture its joint dependence with the other rail operating companies alone while in practice it is also a leading partner in the formation of air-rail alliances (Star Alliance, 2022; Wandelt & Sun, 2022). The dependence measured here is accordingly a lower bound on Deutsche Bahn's role in the European intermodal network, and the whole-system perspective required of these operators extends across both modes.

Second, the form the disruption takes differs with the type of edge withdrawn, and the emergency response should differ with it. Under the DB withdrawal and the service-aware failures that strike the same within-cluster corridors, the network responds by complete disconnection, with the delayed-but-within-tolerance share at 1.2 per cent for DB and below 5 per cent along the service-aware ordering. Under failures of the cross-cluster cut-edges the delayed band roughly doubles, reaching 9.7 per cent of OD-departure buckets at $x = 0.05$ against 4.6 per cent under the service-aware ordering, the partial-degradation case of Cats and Jenelius (2018). Emergency-timetable response should therefore reserve capacity for absorbing complete trip cancellations when the German corridors fail, and

shift toward rerouting through intermodal interlining and air capacity when the international cut-edges fail, where a genuine delayed band exists to manage.

Two methodological conclusions follow, both bearing directly on sub-question 2 and on the comparison between service-aware and topology-based robustness assessments. First, a topology-based assessment misses what determines robustness in two ways. Its edge ordering is statistically indistinguishable from a uniform random benchmark on this network (Section 5.1.2), so it misidentifies the service-critical edges. It also carries no record of how service is organised on the graph, neither the distribution of supply edges across operators, which assigns more than a fifth of the intermodal network to Deutsche Bahn alone, nor the cooperation rules that govern which operators a passenger may combine, and the experiments of this section show both to be highly impactful for robustness. For the regulators and operators who commission robustness assessments, the implication is practical. A structural ranking remains the only option where service information is unavailable and retains its value as a first screening, but where timetable and operator data exist the assessment should incorporate them, tailoring the generic instruments of network science to the operational content of the system under assessment. Second, beyond the small set of corridors that Section 5.1 identified as carrying the heaviest itinerary load, the operators and the rules under which passengers may combine them within one itinerary also determine robustness (Section 5.2.3). The B1-to-B5 cooperation gradient exceeds the service-aware-versus-random gap by a factor of 1.53× (Section 5.2.3). The finding gives empirical form to the market-dynamics challenge that Cats (2025) identifies for the European long-distance market, in which coalition formation among operators is a central open question, and to the fragmentation of the open-access market that Beria et al. (2023) document. The cooperation set C should therefore be treated as a parameter in its own right in future European-network robustness studies, varied alongside the disruption rule. A comparison that fixes C tacitly will misstate the network's robustness envelope by more than the infrastructural variation under study.

The three implications converge on the same corridor and together form the basis on which Chapter 6 answers sub-question 4. The protection instrument is directed at the corridor that the service-aware ranking and the DB withdrawal identify, the rebooking instruments apply to the operators crossing that corridor, and the reframing of cooperation as a parameter in its own right allows future studies to compare robustness across cooperation arrangements rather than across infrastructure alone.

5.3. Type III: macro-scenario disruption

Type III scenarios translate a documented or plausible hazard into the route-and-flight set it disables, and together with the operator-side experiments of Section 5.2 they complete the chapter's answer to sub-question 3 on the extent to which rail-air co-existence enhances robustness across both modes. Two scenario families are reported. Both events are introduced as documented hazards on the network in Section 4.2.4. The ash scenario tests rail-as-substitute-for-air and the flood scenario tests air-as-substitute-for-rail, the mode-symmetric pairing that the rail-centric cooperation-tier design of Section 5.2 could not provide.

This section deepens the structural-versus-service tension established in sections 5.1 and 5.2. The headline empirical finding is that neither topological redundancy nor cross-

modal substitution measurably eases the constraint on the network's robustness under a real macro-shock. Under the ash scenario the intermodal pool retains $R(x) = 0.875$ at the peak ash polygon, but this retention comes from a baseline-immunity effect, that the rail-only baselines never traverse airports, rather than from cross-modal rescue. The cross-modal substitutes that a topology-based perspective would credit as available are filtered out by the service-level reroute tolerance of equation 3.39. The Bernd flood produces the symmetric finding on the rail side, with the air-as-substitute mechanism similarly filtered out by the tolerance bands. Subsections 5.3.1 and 5.3.2 report each scenario in turn, and the policy implications for sub-question 4 are developed in Chapter 6 via the chapter-wide synthesis of Section 6.1.

5.3.1. III.A Volcanic-ash airspace closure (Iceland 2010)

The Iceland 2010 ash scenario evaluates the Type III disruption rule under a real-world air-only hazard whose extent does not correlate with service-aware betweenness.

The scenario operationalises the event through five day-resolved ash-zone polygons reconstructed in EPSG:4326 from the daily ash advisories of the London Volcanic Ash Advisory Centre (London Volcanic Ash Advisory Centre, UK Met Office, 2010) and cross-validated against the per-country closure tallies reported by Eurocontrol (2010). The procedural background of the VAAC advisory chain and the FL000-200 / FL200-350 / FL350-550 ash-concentration banding adopted by EUROCONTROL during the crisis is reviewed by Brooker (2010). The spatial-coherence property of the ash zone, that the affected airspace forms a contiguous geographic block rather than a uniformly random sample of European airspace, is the same property Wilkinson et al. (2012) identify as the source of the disproportionate operational damage on the European air-traffic network. The five named zones (apr15-expand, apr17-peak, apr18-peak, apr19-retract, apr21-residual) cover the documented closure extent on the most-disrupted days of the event. The peak-day scenario closes the 153 airports inside the apr18-peak zone; the union scenario closes the 178 airports that fell inside any zone at any point during the eruption. Each airport closure cascades through both layers of the functional MLTN of Section 3.1.4: every supply-layer flight service $e \in E^{\text{sup}}$ with at least one endpoint in the closed set is invalidated, and every structural air edge incident to a closed airport is removed from the topological layer. The closed-airport set covers 77.7% of the 197 active airports of the intermodal network.

The high $R(x)$ at the peak day is a baseline-composition effect rather than a cross-modal rescue. The 38.95% of OD-departure baselines that are rail-only never traverse airports and are therefore structurally immune to an air-only hazard, and this share alone places a floor of $R(x) \geq 0.39$ on the intermodal retained efficiency under any ash-class event. The true air-to-rail cross-modal rescue counter, defined as the count of OD-departure buckets whose air baseline is replaced by a within-tolerance rail substitute under equation 3.39, is flat zero across both peak-day and union variants. The 1,318 cross-mode counter reported under the peak-day scenario is intra-FUA airport redundancy, a rail-air-rail baseline that finds a parallel rail-air-rail itinerary through a different non-ash airport in the same functional urban area, and is therefore not a mode shift.

The ash scenario is the macro-shock analogue of random failure on the air layer. The peak-day zone spans the full European air geography (Iberia, France, the United Kingdom, the Benelux, Germany, Scandinavia, Italy, and Central Europe) without concentrating on

any subset of high-itinerary-load corridors, and the additional airports added by the apr15-expand and apr21-residual perimeters of the union scenario are operationally peripheral. Two implications for sub-question 4 follow. First, no protection budget concentrated on a small subset of airports can blunt an ash-class event, because the protection-of-specific-airports problem is structurally unsolvable when the disruption covers the majority of the European airport set. Second, the air-to-rail cross-modal rescue mechanism that an infrastructural perspective would credit as the natural absorber is empirically absent under a continent-scale airspace closure, and the robustness lever is therefore rail-side surge capacity rather than cross-modal substitution. A future eruption of comparable scope would require investment in rail-side capacity to absorb the rail-only travellers whose itinerary is now their only mode option, rather than in protecting any specific subset of airports or in expecting an air-to-rail rescue that did not materialise in 2010 and does not materialise in this study's reconstruction.

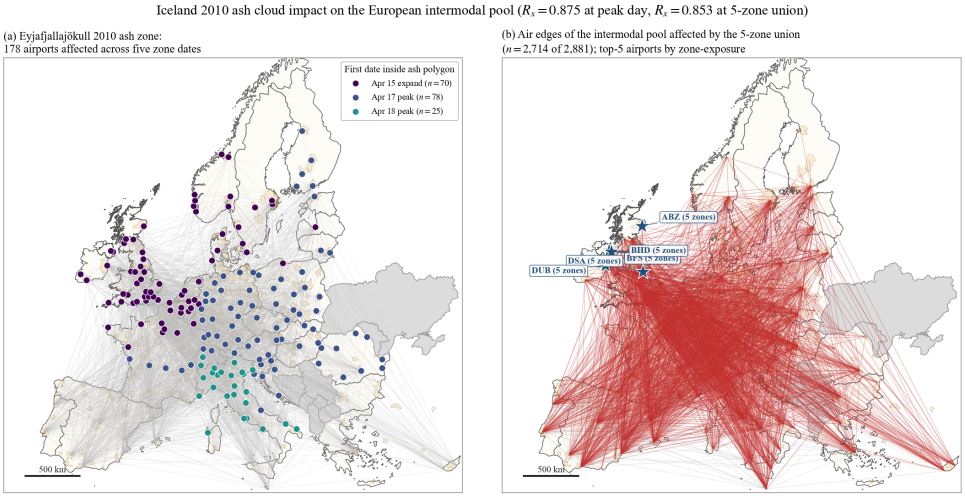

Figure 5.10: Iceland 2010 ash cloud impact on the European intermodal network. Panel (a) reports the 178 affected airports across the five named ash-zone dates (apr15-expand, apr17-peak, apr18-peak, apr19-retract, apr21-residual), colour-coded on a viridis gradient by the earliest zone date on which the airport fell inside the London VAAC ash-zone polygon. Panel (b) reports the air edges of the intermodal network with at least one endpoint inside the five-zone union (drawn in red, n = 2,714 of 2,881 air edges) and the surviving air edges (light grey, n = 167). The top-5 airports by zone-exposure (number of distinct zones the airport fell inside) are starred. Affected-airport count and air-edge impact: ash_peak (apr18, 153 airports closed) delivers $R(x) = 0.875$ with $f_{\text{disc}} = 20.6\%$ and $\Delta\bar{D} = 37.1$ minutes; ash_union (five-zone union, 178 airports closed) delivers $R(x) = 0.853$ with $f_{\text{disc}} = 24.5\%$ and $\Delta\bar{D} = 44.1$ minutes. The ash zone covers 94 per cent of the air edges of the intermodal network yet retains 87.5 per cent of the intermodal travel utility at the peak day, a baseline composition effect rather than cross-modal rescue: 38.95 per cent of the OD-departure baselines are rail-only and never traverse airports, so the ash zone is structurally inapplicable to them by mode-specific baseline immunity. The true air-to-rail cross-modal rescue under the ash zone is flat zero; the 1,318 cross-mode bucket counter recorded under the peak-day scenario is intra-FUA airport redundancy, a rail-air-rail baseline that finds a parallel rail-air-rail itinerary through a different non-ash airport in the same functional urban area, not a mode-shift onto rail.

5.3.2. III.B Continental flood (river-basin scenarios and Bernd 2021)

The continental flood scenario uses two complementary framings. The river-basin variant applies the Type III disruption rule to the union of the Danube middle, Rhône, Po, Vistula, and the extended Rhine basin padded north to Rotterdam and Düsseldorf, together with the union of all five basins. The event-based comparator uses the Copernicus Emergency Management Service rapid-mapping for the July 2021 Bernd storm, EMSR517 for Germany, EMSR518 for Belgium, and EMSR520 for the Netherlands (Copernicus Emergency Management Service, 2021a, 2021b, 2021c). The event-based variant is reported at both the Europe-wide bucket scope and at the intra-country scope to remove the dilution from long-haul OD windows whose itineraries never traverse the flooded region.

Figure 5.11 maps the Copernicus EMS flood polygons onto the rail layer of the structural MLTN for the three national mappings. The observed-event polygon is drawn in red hatching and the 5 km hydrological buffer applied to account for indirect damage to embankments, track substructures, and signalling cabinets is drawn in light blue. The surviving rail edges are drawn in grey and the invalidated rail edges in red (inside the buffer) and orange (outside the buffer but on a rail service whose remaining edges include at least one inside the buffer, so the service cannot operate end-to-end). Stations whose geographic location falls inside the 5 km hydrological buffer are marked as red circles. The German EMSR517 panel shows the densest invalidation footprint, with 77 such stations inside the buffer and 365 rail edges removed primarily along the Ahr valley, the Eifel line, and the Erftstadt and Cologne-Trier catchments. The Belgian EMSR518 panel shows 25 affected edges on the Verviers-Eupen axis, and the Dutch EMSR520 panel shows 16 affected edges on the Maaslijn.

The 2021 Bernd flooding shows where a real macro-shock falls on the geographic-correlation dimension of Section 5.1, between the random draw and the service-aware worst case. The actual event is neither. Its 365 rail-edge removals do less damage than the same number of edges taken from the top of the service-aware ranking, and the damage multiplier of 1.69× on the intra-German bucket scope measures how much worse that worst case would have been. Table 5.6 shows the same dependence on placement at the national scale. Intra-German and intra-Dutch journeys experience an 8.1× contrast in disconnection share under the same Atlantic weather system, because the German invalidated edges are on the main lines of the European long-distance network while the Dutch invalidated edges are on the Maaslijn, whose passengers reach their destinations through the IJsselmond and Limburg corridors that the flood does not reach. Bernd is therefore the chapter's real-event measurement of how strongly the damage of a hazard depends on where it falls, and the policy implications follow from its comparison with the ash scenario in Section 5.3.3.

5.3.3. Policy insights from the ash and flood scenarios

The ash and Bernd scenarios are the two best-documented European macro-shocks of opposite spatial signature. The air-side hazard reaches 77.7% of operating airports at once, while the rail-side hazard removes 365 structural edges in one spatially clustered region. Read together through the four-state classifier of equation 3.40, they give the chapter's clearest evidence that topology-based assessments overestimate service-level robustness under macro-shocks. Under ash the intermodal pool retains $R(x) = 0.875$, and the effec-

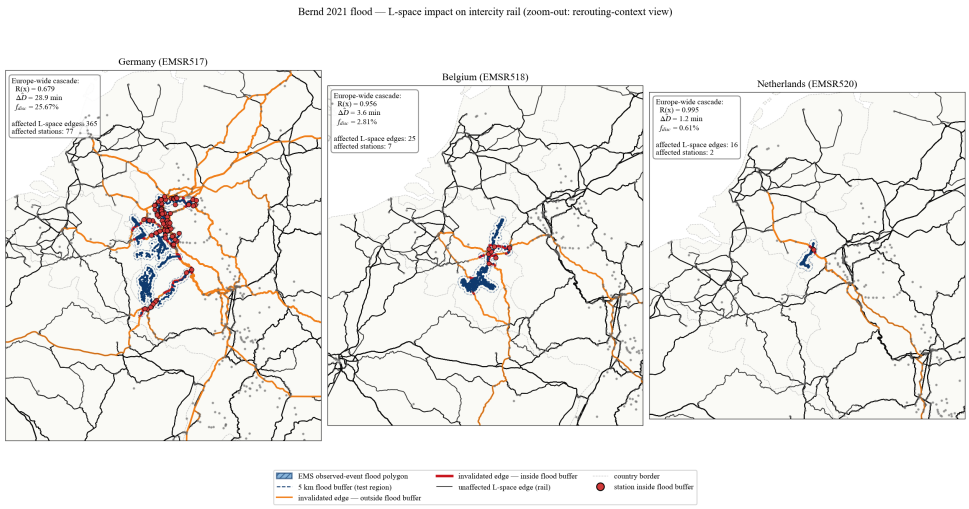

Figure 5.11: Bernd 2021 flood, L-space rail impact for the three Copernicus Emergency Management Service mappings (EMSR517 Germany, EMSR518 Belgium, EMSR520 Netherlands). Affected stations are marked as red circles inside the 5 kilometre flood buffer (light blue). The EMS observed event flood polygon is in red hatching. L-space rail edges are drawn in grey; invalidated edges are in red (inside the buffer) or orange (outside the buffer but on a route that traverses the buffer). The cumulative Europe-wide cascade outcome is annotated in the upper-left of each panel, $R(x) = 0.679$ at EMSR517 (DE), $R(x) = 0.956$ at EMSR518 (BE), and $R(x) = 0.995$ at EMSR520 (NL). The Cologne to Liège to Brussels triangle emerges as the single most vulnerable sub-corridor under any of the three mappings, with 94 of the 365 edge invalidations under the German mapping concentrated on the Cologne station cluster.

Table 5.6: Bernd 2021 flood (Copernicus EMS) intra-country retained efficiency. Each row reports one EMS mapping applied to the intermodal pool at the listed bucket scope. The intra-country scope restricts the bucket set to OD-departure combinations whose origin and destination both lie in the listed country, so that the headline $R(x)$ does not absorb the dilution from long-haul OD pairs that never traverse the flooded region.

EMS mapping	Edges	Scope	$R(x)$	f_{disc} (%)	$\Delta\bar{D}$ (min)
EMSR517 (DE)	365	Europe-wide	0.679	25.7	28.9
EMSR517 (DE)	365	intra-DE	0.510	47.8	37.6
EMSR518 (BE)	25	intra-BE	0.767	26.2	19.1
EMSR520 (NL)	16	intra-NL	0.944	5.9	6.7

tively disconnected fraction equals the strictly disconnected fraction ($f_{\text{eff-disc}} = f_{\text{disc}} = 0.206$). Every disrupted bucket has lost every substitute, and the gap between the two fractions is empty by construction, because the 38.95% of OD-departure baselines that are rail-only are structurally immune to an air-only disruption. Under Bernd the pool drops to $R(x) = 0.510$ on the intra-German scope and the same gap opens, with $f_{\text{eff-disc}} = 0.504$ against $f_{\text{disc}} = 0.478$. The difference is approximately 81,000 OD-departure buckets in which a structurally feasible cross-modal substitute exists but arrives outside the reroute tolerance of equation 3.39. This 81,000-bucket gap is the chapter's bucket-level measurement of the divergence that Cats (2025) identifies as a central methodological challenge for European intermodal robustness assessment, and it shows at the macro-shock scale

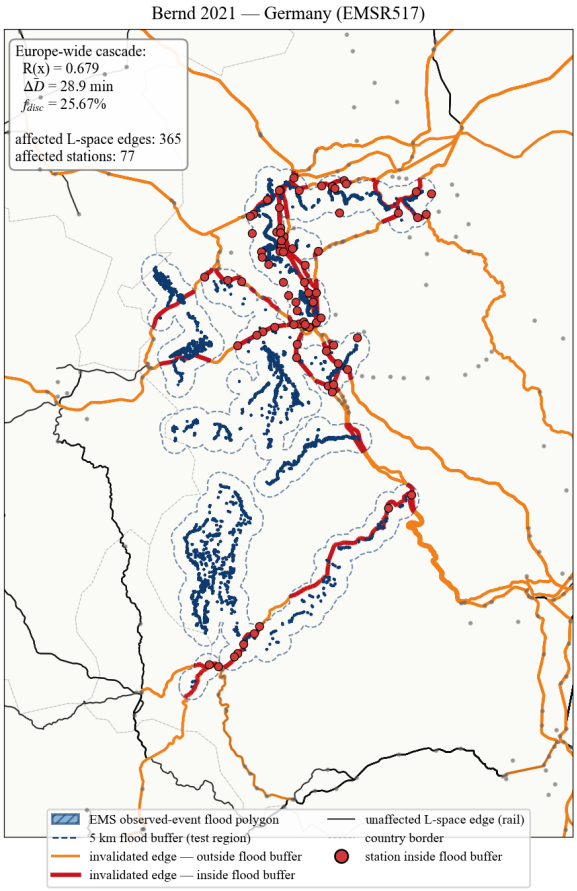

Figure 5.12: Bernd 2021 flood, Germany-only close-up of the L-space rail impact under the Copernicus EMSR517 mapping. The 365 invalidated rail edges concentrate on the Ahr valley, the Eifel line, and the Erftstadt and Cologne-Trier catchments. The 77 affected stations are marked as red circles inside the 5 kilometre flood buffer. The affected area is adjacent to but does not coincide with the German corridors elevated by the service-aware ranking (Section 5.1). The Rhine-Ruhr industrial corridor and the Frankfurt-Mannheim main line are geographically close to the Bernd zone but were spared by the actual rainfall distribution.

the distinction between structural connectivity and journey continuity at usable service quality.

The two scenarios show the protection question from opposite sides, and what helps differs by mode. The ash hazard reaches over three-quarters of the operating airports at once, so no protection of a small set of airports or air corridors can reduce the damage. The investment that helps against an air-side hazard is therefore rail capacity to absorb the stranded passengers, as Section 5.3.1 concludes, and the rail-only baselines already withstand a large share of an air-side hazard without any protection spending. The Bernd hazard strikes one small region instead, but its damage depends on which part of the German cluster the next flood reaches, the margin that the damage multiplier of Section 5.3.2 mea-

sures. Protecting the heavily loaded German corridors is therefore effective only where the flood-prone parts of the cluster are known in advance. European robustness planning should accordingly treat the two modes separately, rail capacity against air-side hazards and corridor protection guided by flood forecasting against rail-side hazards, rather than fund both from a single protection budget.

Figure 5.13 places the two macro-shocks against the random and service-aware benchmarks of Section 5.1 on the same three indicators. Ash falls outside the targeting curves entirely, because its closure covers nearly the whole air network at once rather than any chosen set of edges. No within-band cross-modal rescue is available under it, and the retained efficiency survives only through the rail-only share of the baseline composition (Section 5.3.1). Bernd falls between the random benchmark and the service-aware worst case at matched edge count. The actual EMSR517 mapping reaches $R(x) = 0.679$ Europe-wide, while a removal of the same number of edges taken from the top of the service-aware ranking would reach $R(x) \approx 0.17$, the distance quantified in Section 5.3.2. The figure therefore extends the geographic-correlation spectrum that Mattsson and Jenelius (2015) formalises from synthetic edge failures to documented macro-shocks, and it confirms that the robustness dimension the chapter identifies at the edge-targeting scale governs real-world hazards as well.

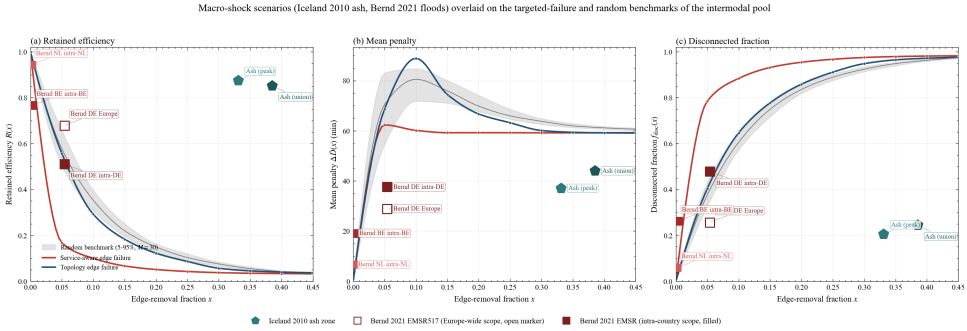

Figure 5.13: Macro-shock scenarios (Iceland 2010 ash plume, Bernd 2021 floods) overlaid on the targeted-failure and random benchmarks of Section 5.1 on the intermodal pool of 6,760 edges. Panels report (a) retained efficiency $R(x)$, (b) mean penalty $\Delta\bar{D}(x)$ in minutes, and (c) disconnected fraction $f_{\text{disc}}(x)$ against the edge-removal fraction x. Background lines reproduce the service-aware ($\hat{b}$) and topology-based (b^{topo}) edge-failure curves with the random-failure 5 to 95% envelope ($M = 30$ seeds) of Figure 5.2. Teal pentagons mark the ash scenarios at the peak and union zones. Brown squares mark the Bernd EMSR scenarios, with the open square for the Europe-wide bucket scope and the filled squares for the intra-country bucket scopes; EMSR517, EMSR518, and EMSR520 are positioned at $x = 365/6{,}760$, $25/6{,}760$, and $16/6{,}760$ respectively, and the Europe-wide and intra-DE rows of EMSR517 share the same x because they describe the same 365-edge removal under two bucket-scope masks.

5.4. Sensitivity analysis

This section is a methodological annex placed before the synthesis of Section 6.1. Before tying the findings of sections 5.1 through 5.3 into a single causal chain, we confirm that the headline retained-efficiency numbers and the ordering of stations, edges, and corridors are robust to the parameters of the pipeline that a reviewer is most likely to question. Four parameters are perturbed within their plausible ranges, namely the tolerance band

of equation 3.39, the random-failure Monte-Carlo seed count M, the path-enumeration $W^{\text{rail}}_{\max}$ cap of Section 3.2.3, and the two Wardman generalised-travel-cost multipliers μ_{wait} and ν_{move} of Wardman et al. (2016). The minimum-transfer-gap floor of Section 3.2.3 is bounded analytically against the cooperation-tier evidence of Section 5.2. Figure 5.16 reports the four empirical sweeps in panels (a) to (e) and the analytical tornado-bound across all dimensions in panel (f). The headline finding of this section is that no parameter inside its plausible range can close the chapter's central retained-efficiency gap of 0.389 units between the service-aware and topology failure curves at $x = 0.05$. The chapter's topology-versus-service split is therefore not an artefact of parameter tuning.

5.4.1. Sensitivity to the airport boarding ceiling $W^{\text{air}}_{\max}$

Section 3.2.3 set the airport boarding ceiling $W^{\text{air}}_{\max}$ loosely, at 24 hours, to avoid pruning feasible connections in the sparse air layer. This section tests how much that choice matters. The waiting time an itinerary spends at an airport between legs, comprising the mandatory boarding floor and any further layover, is dominated by the structural floors rather than by idle time. A pure air-to-air change has a median airport wait of 185 minutes, a rail-to-air change a median of 223 minutes, and even the cost-minimising air-involving itineraries retained per bucket carry a median of 310 minutes, since the boarding floors of 90 to 120 minutes cannot be avoided.

The ceiling itself does not bind. No enumerated air itinerary reaches it, the longest observed airport wait being 1,240 minutes. Long waits are nonetheless a real feature of the sparse air layer, as Table 5.7 shows. Lowering the ceiling to twelve hours would remove every air option from only 0.5 per cent of origin destination departure buckets, whereas lowering it to six hours would do so for 22 per cent. The relations that depend on a long wait are therefore the thinly served ones, for which the single-month air schedule offers no quicker feasible connection, so the loose ceiling preserves their air option. A behaviourally tighter ceiling, which would reclassify these relations as rail-only, is reported here as a bound rather than adopted.

Table 5.7: Effect of the airport waiting ceiling on the enumerated air itineraries. The second column gives the share of air-involving alternatives whose aggregate airport waiting time exceeds each value. The third column gives the share of origin destination departure buckets that lose every air-involving alternative once the ceiling is set to that value.

Ceiling	Air alternatives above it	Buckets losing all air
6 hours	36.3%	22.1%
8 hours	7.2%	4.7%
12 hours	0.7%	0.5%
18 hours	0.01%	0.003%
24 hours	0%	0%

5.4.2. Tolerance band of equation 3.39

The tolerance band is the most directly testable parameter of the chapter and the one a reviewer is most likely to question. The canonical band is 60 minutes for rail-only baselines, aligned with the one-hour delay-compensation threshold of (European Parliament and

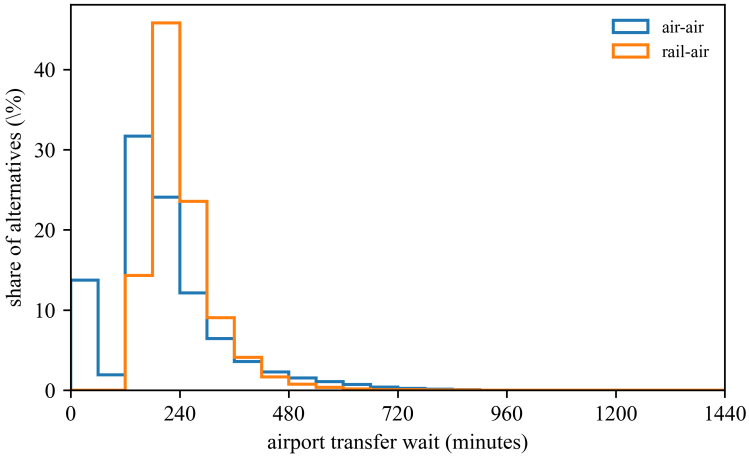

Figure 5.14: Distribution of the aggregate airport waiting time for air-to-air and rail-to-air itineraries. The mass sits well below the 24-hour ceiling, which no enumerated itinerary reaches.

Council of the European Union, 2021), and 180 minutes for air-involving baselines, aligned with the three-hour scale of (European Parliament and Council of the European Union, 2004). Panel (a) of Figure 5.16 sweeps the rail-side band T_{rail} over $\{60, 90, 120, 180, 240\}$ minutes while preserving the $1:3$ rail-to-air ratio that the regulatory anchor imposes, and reports the within-tolerance share of cross-modal substitutes for the four cluster-shaped scenarios of the chapter, namely the service-aware failure, the topology failure, the Rhine-extended counterfactual, and the Bernd EMSR517 flood. Under the service-aware failure the within-tolerance share rises from 0.30 per cent at the canonical $T_{\text{rail}} = 60$ to 0.93 per cent at $T_{\text{rail}} = 240$, namely a fourfold relaxation of the band moves only an additional sub-one-per-cent of substitutes into the within-tolerance class. Panel (b) reports the corresponding disconnected populations on a logarithmic axis, and the four lines are visually flat across the sweep range. The chapter's headline finding that cross-modal substitution provides paths but not service quality therefore does not depend on the specific tolerance numbers anchored by the two passenger-rights regulations.

5.4.3. Random-seed Monte-Carlo convergence

The random benchmark of Section 5.1.2 uses $M = 30$ Monte-Carlo seeds at each intensity grade. The count is set by the size of the effect it must resolve. The seeds are independent edge-removal permutations, so the random benchmark mean is an ordinary sample mean whose Monte-Carlo standard error shrinks with the seed count as $\sigma_R(x)/\sqrt{M}$ (Law & Kelton, 2000). At the diagnostic intensity $x = 0.05$ the seed-to-seed standard deviation is $\sigma_R \approx 0.048$, so the standard error of the benchmark mean at $M = 30$ is only about 0.009. The service-aware failure lies 8.5 standard deviations below that mean, far beyond this sampling error. Figure 5.15 confirms each step. The benchmark mean converges on its full-pool value of 0.580 by $M = 30$ in panel (a). The thirty random values are close to normal in panel (c), which supports the Gaussian z-scores and p-values reported in Sec-

tion 5.1.2. The separation stays beyond six standard deviations even at $M = 15$ in panel (b). The seed-to-seed dispersion is largest at $x = 0.05$ and declines toward $x = 1$ in panel (d), so the count is tested at its most demanding intensity and holds across the whole removal range. A larger seed count would only tighten the benchmark, so the random-versus-service-aware z-score of -8.51 is not a Monte-Carlo artefact.

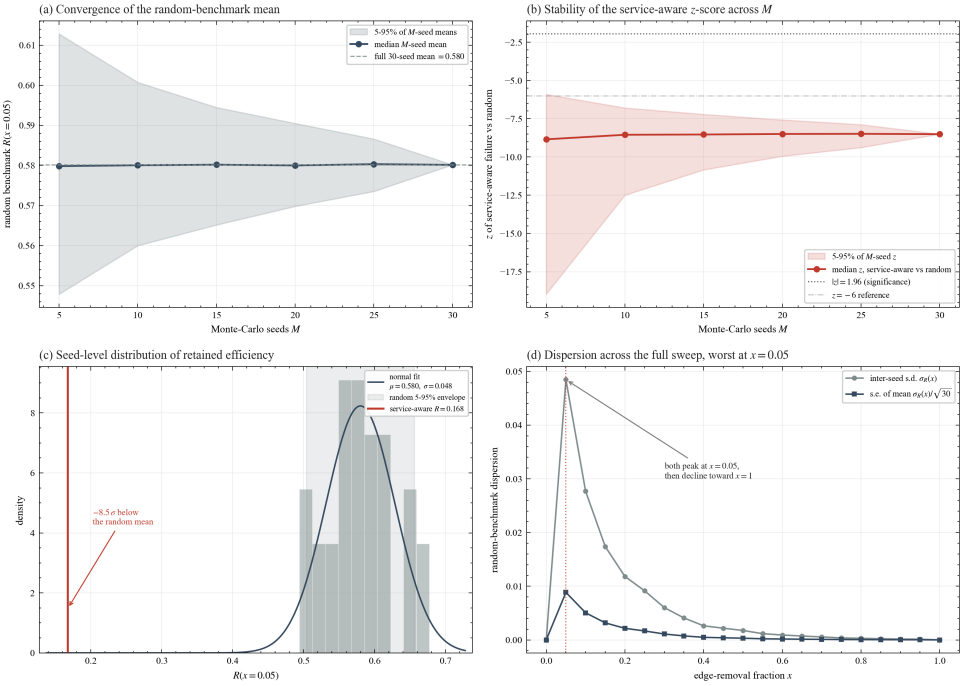

Figure 5.15: The $M = 30$ random benchmark is sufficient across the entire disruption sweep, and the diagnostic intensity $x = 0.05$ is its most demanding point rather than a favourable one. Panels (a) to (c) examine $x = 0.05$ on the intermodal pool. Panel (a) subsamples the thirty seeds at $M \in \{5, 10, 15, 20, 25, 30\}$, and the benchmark mean converges on the full-pool value of 0.580 as the 5 to 95 per cent range of the M-seed means narrows with increasing M. Panel (b) presents the z-score of the service-aware failure against the random benchmark estimated from M seeds. It remains close to -8.5 and exceeds -6 in magnitude even in the least extreme subsample at $M = 15$, far below the $|z| = 1.96$ significance threshold. Panel (c) presents the thirty random retained-efficiency values with their normal fit ($\mu = 0.580$, $\sigma = 0.048$), and the service-aware value of 0.168 falls 8.5 standard deviations below the random mean. Panel (d) presents the inter-seed standard deviation and the Monte-Carlo standard error of the random-benchmark mean across the full removal range. Both are largest at $x = 0.05$ and fall monotonically toward $x = 1$, so the seed count is tested most stringently at the diagnostic intensity and is over-satisfied at every higher intensity.

5.4.4. Same-station boarding ceiling $W^{\text{rail}}_{\max}$

The $W^{\text{rail}}_{\max}$ parameter caps the same-station waiting time tolerated during the path-enumeration construction. The canonical pipeline uses $W^{\text{rail}}_{\max} = 60$ minutes, consistent with the median observed transfer dwell of 22.6 minutes in the path pool and within the European connecting-time guidance of (International Air Transport Association, 2024).

Panel (d) of Figure 5.16 sweeps $W^{\text{rail}}_{\max}$ over $\{15, 30, 45, 60, 90\}$ minutes and reports two saturation curves. The zero-transfer pool is invariant by construction. The one-transfer pool grows from 2.59 million paths at $W^{\text{rail}}_{\max} = 15$ to 5.21 million at $W^{\text{rail}}_{\max} = 90$, with the marginal-gain curve flattening at $W^{\text{rail}}_{\max} = 60$, namely the path pool gains only an additional 7.5 per cent of paths between 60 and 90 minutes against the 47 per cent contraction between 60 and 15 minutes. The one-transfer OD-pair coverage grows from 112,873 to 169,948 pairs across the same range and saturates similarly. The canonical choice therefore sits past the elbow of the marginal-gain curve, namely tighter values remove feasible itineraries and looser values yield diminishing returns. The $R(x)$ sensitivity to $W^{\text{rail}}_{\max}$ is bounded above by the path-pool composition shift and is of order 0.05 retained-efficiency units, well below the chapter's headline service-aware-versus-topology gap.

5.4.5. Generalised-travel-cost multipliers

The thesis GTC formula combines in-vehicle time with waiting time through two multiplicative weights $\nu_{\text{move}} = 1.87$ and $\mu_{\text{wait}} = 1.76$ of Wardman et al. (2016). The enumerated path set is invariant under any sweep of these weights because the upstream filter operates on travel time rather than on GTC, namely the multipliers change how each path's GTC is ranked but not which paths exist. Panel (e) of Figure 5.16 sweeps μ_{wait} across $[1.70, 2.00]$ in steps of 0.05 at the fixed $\nu_{\text{move}} = 1.87$, bounded below by the European long-distance value-of-time interval of Wardman et al. (2016) and above by the upper bound of the UK Department for Transport TAG Unit A1.3 (Department for Transport, 2024). The mean GTC at the one-transfer level rises by 2.45 per cent (from 261.87 to 268.30 minutes), the median by 2.92 per cent (from 233.20 to 240.00 minutes), and the 95th percentile by 1.95 per cent across the full range. The retained-efficiency ratio $R(x) = E_G(x)/E_G(0)$ is the ratio of two GTC-weighted efficiencies that scale together under a proportional multiplier shift, and this 2 to 3 per cent absolute shift cancels in the ratio. The robustness ordering of stations, edges, and corridors therefore does not depend on the specific Wardman value used.

5.4.6. Headline-bound tornado

Panel (f) of Figure 5.16 reports the upper bound on $|\Delta R(x = 0.05)|$ that each sensitivity dimension could in principle deliver, ranked from largest to smallest. The chapter's headline service-aware-versus-topology gap of $0.557 - 0.168 = 0.389$ retained-efficiency units is drawn as a vertical reference. Every individually-bounded dimension sits below this gap. The minimum-transfer-gap floor of Section 3.2.3 is bounded analytically by the cooperation-tier B1-to-B3 jump of 0.255 retained-efficiency units reported in Section 5.2, namely an upper bound on the institutional interline-feasibility gap of which the minimum-gap floor is a strict subset, and is still 0.134 units below the headline gap. The $W^{\text{rail}}_{\max}$ empirical bound at 0.050 units is 0.339 units below the headline gap. The remaining dimensions (the tolerance band, the random-seed count, and the Wardman multipliers) deliver bounds that are negligible on the same scale. The chapter's central finding that topology- and service-based assessments diverge systematically is therefore not an artefact of any single pipeline parameter inside the testable range.

5.4.7. Summary

Four pipeline parameters are tested empirically and one further parameter is bounded analytically against the chapter's cooperation- tier evidence. The tolerance band of equation 3.39 is robust to a fourfold relaxation. The two Wardman multipliers of Wardman et al. (2016) are analytically invariant in the retained-efficiency ratio. The random benchmark converges at $M = 30$ with a Monte-Carlo standard error of approximately 0.0009 at $x = 0.05$. The path-enumeration $W_{\max}^{\text{rail}}$ cap and the analytically bounded minimum-transfer-gap floor both sit below the chapter's headline service-aware-versus-topology gap of 0.389 retained-efficiency units. The robustness ordering of European edges and corridors reported in sections 5.1 through 5.3 is therefore not an artefact of parameter tuning.

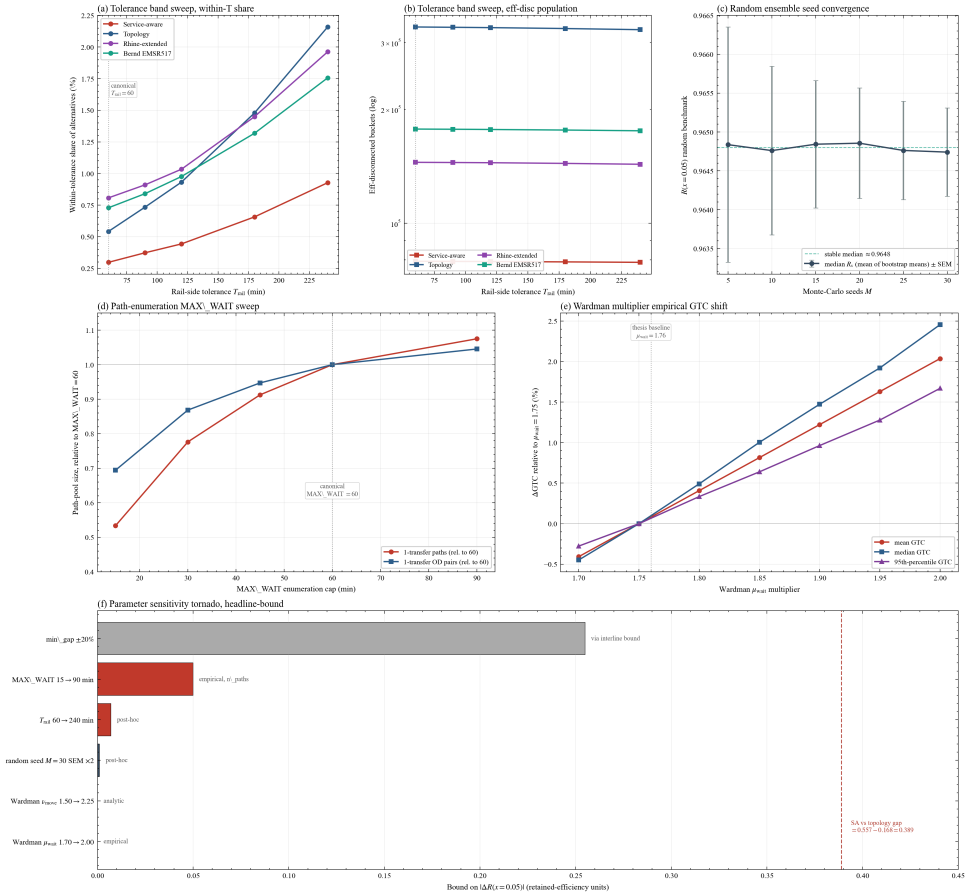

Figure 5.16: Canonical sensitivity analysis on the band-180 three-year pipeline. Panel (a) sweeps the rail-only tolerance $T_{\text{rail}} \in \{60, 90, 120, 180, 240\}$ minutes (with the 1 : 3 rail-to-air ratio preserved) and reports the within-tolerance share of cross-modal substitutes per cluster-shaped scenario, namely the service-aware failure, the topology failure, the Rhine-extended counterfactual, and the Bernd EMSR517 flood. Panel (b) reports the corresponding disconnected populations on a logarithmic axis; all four lines are essentially flat across the sweep. Panel (c) reports the Monte-Carlo seed convergence on the random benchmark, namely the median of bootstrap means at $x = 0.05$ stabilises and the standard error falls as $1/\sqrt{M}$. Panel (d) reports the path-enumeration $W^{\text{rail}}_{\max}$ sweep over $\{15, 30, 45, 60, 90\}$ minutes; the one-transfer path pool and OD-pair coverage saturate at the canonical $W^{\text{rail}}_{\max} = 60$. Panel (e) reports the empirical Wardman-multiplier sweep at the one-transfer level over $\mu_{\text{wait}} \in [1.70, 2.00]$; the absolute GTC shift is at most 2.92 per cent and cancels in the retained-efficiency ratio. Panel (f) reports the upper bound on $|\Delta R(x = 0.05)|$ that each sensitivity dimension could in principle deliver, ranked against the chapter's headline service-aware-versus-topology gap of 0.389 retained-efficiency units drawn as a vertical reference. Every individually-bounded dimension sits below the headline gap.

6

Conclusion

6.1. Cross-scenario synthesis and answers to the research questions

The three experiment families of Chapter 3.4 are not designed to stand alone. Type I separated the topological and the service-level perspective on the same graph, type II traced the divergence to its operator-side carrier, and type III tested the two mechanisms that might still have rescued the network. This closing section draws them together into a single line of reasoning and follows it back to the questions the thesis set out to answer. The reasoning proceeds by elimination, ruling out one explanation at a time, and it returns at the end to the dual representation it began with, since each step resolves one of those questions and the steps are only meaningful in sequence.

6.1.1. Synthesis of the three experiments

The reasoning begins with the divergence that the dual network makes visible. The unweighted topology ranking correctly locates the structural cut-edges that hold the British Isles, Iberia, Scandinavia, and the continental core attached, the connectivity question answered by the percolation tradition of Albert et al. (2000), the public-transport benchmark of von Ferber et al. (2009), and the textbook account of Barabási (2016). It does not locate the edges that carry passenger throughput. At the removal fractions most relevant to planning the topology ordering is statistically indistinguishable from a uniform-random allocation, while the service-aware ordering departs from it sharply (Section 5.1.2), so a protection budget spent by topology betweenness buys no more retained efficiency than the same budget spent at random. Incorporating service-level information does not rescale the robustness estimate, it relocates what counts as critical, the continental-scale form of the divergence that Pien et al. (2015) measured between unweighted betweenness and air-traffic delay, that Zhou et al. (2019) confirmed on the weighted air network, and that Cats (2025) names among the open challenges for European long-distance network science.

This relocation is not diffuse. It has a definite geography, and it is the same geography wherever the chapter looks. The service-aware edge ranking elevates the Rhine-Ruhr cor-

ridor, the Frankfurt-Mannheim main line, and the South German high-speed alignments, the Deutsche Bahn withdrawal removes the same corridors as a single operator's service area, and the July 2021 Bernd flood reaches its worst intra-German damage on the same interior (sections 5.1.1, 5.2.1, and 5.3.2). A path count, a supply-layer operator cascade, and a curved-geometry polygon intersection from Copernicus emergency mapping share no common computation, yet all three contract onto the same Rhine-Alpine interior that Otsuka et al. (2017) describe as the constrained mixed passenger-freight axis of the European network and that Witlox et al. (2022) identify as the region where high-speed services lose operating speed. Adding the Iceland 2010 ash zone as a fourth, air-side scenario, Figure 6.1 shows the four overlays converging on the same sub-network. The convergence rules out a methodological artefact and marks the German interior as a structural property of the network, the region where the service-level loss of robustness concentrates. The air layer, by contrast, organises its own criticality around the hub feeders into London Heathrow and the spokes through Schiphol and Charles-de-Gaulle that Cardillo et al. (2013) and Pellegrini and Rodriguez (2013) describe, and these edges fall outside the German interior entirely.

With the location of this robustness loss established, the two mechanisms that might still restore robustness are excluded in turn. Structural redundancy is the first, and it is unavailable, since the topology ranking already proved no better than random at the removal fractions that matter. Cross-modal substitution is the second, and the macro-shock scenarios show that it provides paths without journeys. Under the air-only ash hazard the network retains most of its efficiency only because the rail-only baselines never traverse an airport, not because an air-to-rail rescue activates, and under Bernd the buckets that keep a structurally feasible cross-modal substitute receive it outside the tolerance band (sections 5.3.1 and 5.3.2). A structural-feasibility audit counts these substitutes as available, and the service-level criterion of equation 3.39 discards them. This upgrades the binary intermodal-dependency measure ρ_i of Ippolito and Cats (2024) into a service-quality measure, and it carries the partial-degradation framework of Cats and Jenelius (2018) and Cats et al. (2017) to the macro-shock scale, with the practitioner constraints behind the undelivered substitute documented by Toet et al. (2026).

What remains once both rescue mechanisms are excluded is contractual. The cooperation tiers span the distance from a single-operator ticket to full modal interline, and that span exceeds the worst infrastructural gap the chapter measures, as the answer to the main research question below quantifies, so the binding constraint on European long-distance robustness is contractual rather than physical. The same conclusion is reached from the supply side by O'Sullivan and Patel (2004), Beria et al. (2023), and Adler et al. (2010), and the mechanism is simply that cooperation converts alternatives already present in the network into journeys a passenger can book, the missing delivery layer that Toet et al. (2026) describe. The reasoning therefore closes where it opened, on the dual representation, because each step resolves one of the questions the thesis set out to answer, and only together do they show how robust these networks really are.

Table 6.1 records the two robustness indices behind this reasoning, the cumulative robustness-loss integral F and the accumulated delay Q, for the intermodal network and for its rail-only and air-involving constituents under each of the three failures. The service-aware failure is the most damaging on every network. The rail-only constituent loses

retained efficiency the fastest, while the air-involving constituent keeps more of it but pays by far the largest rerouting penalty.

Table 6.1: Cumulative robustness-loss integral $F = \int_0^1 \big(1 - R(x)\big)\,dx$ and accumulated delay $Q = \int_0^1 \Delta\bar{D}(x)\,dx$ (minutes), for the intermodal network and its rail-only and air-involving constituents (the three rows of Figure 5.2), under the service-aware and topology targeted failures and the random benchmark. A larger F is a steeper loss of retained efficiency over the removal range. Values are read from the arrays of Section 5.1.

	service-aware		topology		random	
Pool	F	Q	F	Q	F	Q
intermodal	0.934	58.1	0.891	61.2	0.890	62.0
rail-only constituent	0.965	60.8	0.927	72.5	0.911	73.4
air-involving constituent	0.858	167.2	0.805	158.2	0.838	158.9

6.1.2. Answer to the main research question

To what extent are European air and rail networks robust to disruptions when service-level characteristics are incorporated into their network representation?

For the European air and rail system, robustness is not a fixed property of the network. It depends on how the network is represented, on how severe and how targeted the disruption is, and on the rules under which stranded passengers may be rebooked. Treated as a graph of connections, the system appears robust, since removing even the most central edges does little more damage than removing the same number of edges at random ($z_R = -0.47$ at $x = 0.05$). Treated as the set of scheduled services a passenger can use, it is not, with retained efficiency falling to $R(x) = 0.168$ under the service-aware ordering against 0.557 under the topological one at the same removal fraction. The gap widens as the disruption shifts from random failure to the targeted removal of heavily used edges, and as the rebooking rules tighten.

The experiments point in the same direction. Service-level information does not confirm the robustness that the graph suggests. It moves the critical set away from the thin edges that keep the map connected and onto the smaller set of corridors through which most feasible journeys run. Alternative rail paths rarely restore the loss, because they depend on the same corridors that fail. Air alternatives rarely restore it either, because they arrive too late or at too high a generalised cost to count as usable journeys. Once both forms of redundancy are ruled out, what remains decisive is the booking side.

That booking side is the set of rules deciding whose services a stranded passenger may use. The alternatives a disruption removes largely still exist in the network, but they become usable only as cross-operator and cross-modal cooperation widens. Under a strict single-operator ticket the network is at its least robust, and under full cross-operator and cross-modal interlining the same network recovers almost to its undisrupted state. The span between those two arrangements, 0.624 retained-efficiency units, exceeds the 0.409-unit gap between the worst targeted failure and random failure of the same size by a factor of 1.53. Robustness at the level passengers experience is therefore set less by what has been built than by how freely the built system can be used across operators and modes,

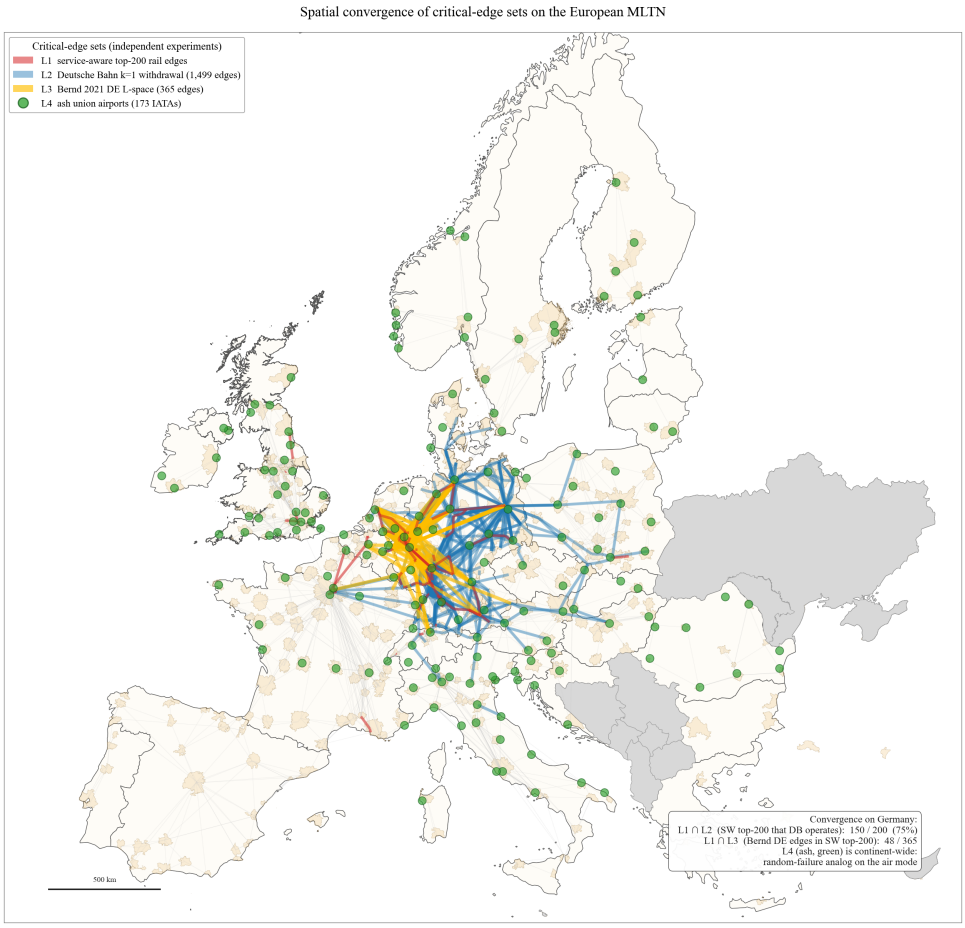

Figure 6.1: Four procedurally independent diagnostics overlaid on one map of the intermodal network. Red marks the top-200 rail edges by service-aware betweenness, blue the supply-layer edges removed by the Deutsche Bahn $k = 1$ withdrawal, yellow the structural rail edges invalidated by the Copernicus EMSR517 mapping of the Bernd 2021 flood, and green circles the airports closed by the Iceland 2010 ash union. The first three overlays concentrate on the same German interior, with 150 of the top-200 service-aware edges operated by Deutsche Bahn and 48 of the 365 Bernd edges among the service-aware top-200, while the ash closure spreads across the whole continent as the air-side counterpart of random failure. The four diagnostics share no common computation, so the convergence is evidence of a structural feature of the network rather than of a single ranking artefact.

and the largest available improvement lies in broadening that cooperation rather than in new infrastructure.

Stated as a single claim, the answer is the following.

> Once service-level characteristics enter the network representation, the European rail and air networks prove far less robust as services than they appear as graphs of topological connections. Their robustness is governed less by

topological connectivity than by the concentration of feasible journeys on a small number of corridors. Ranking edges by pure topological betweenness does not identify those whose removal most damages the journeys passengers can make, measured by generalised travel cost. The edges that topological analysis favours, which hold the graph together, are not those that carry the most feasible journeys. These corridors lie almost entirely within Germany's domestic network, whose service depends heavily on Deutsche Bahn, so that the withdrawal of this single operator alone ranks among the most damaging failures in the study. Cross-modal substitution rarely supplies a usable alternative, whether after the operator withdrawal or under a macro-shock. The robustness passengers experience is therefore set less by what has been built than by the cooperation rules under which they may be rebooked. Broadening cross-operator and cross-modal rebooking recovers much of the robustness the network already contains.

6.1.3. Answer to sub research questions

The answer above is carried by the four sub-questions that structure the thesis, and this section answers each in turn, linking the network representation, the robustness experiments, and the policy analysis back to the chapters that produced them.

Question 1 on network representation - How can the European long-distance air and rail networks be represented as a multilayer graph whose links carry both infrastructure and service-layer information, namely the operating route or flight, the ordered sequence of stations or airports it serves, the scheduled departure and arrival times, and the in-vehicle running time?

The European long-distance air and rail networks can be represented as such a multilayer graph, and Chapter 3 does so as a pair of coupled networks rather than one. Section 3.1.3 fixes the physical skeleton, a mode-layered structure whose supra-adjacency in equation 3.9 keeps rail and air as separate layers joined only by ground-access coupling. Section 3.1.4 overlays the service detail on top of it. Each link in the functional network is a scheduled service rather than a static connection, and through equation 3.12 it carries the operating route or flight, its ordered sequence of stations or airports, the scheduled departure and arrival times, and the in-vehicle running time recorded in equation 3.13, from which the generalised travel cost of every itinerary is built. Infrastructure and service are thus carried together by separating the physical connection from the timetabled service that realises it.

Modelling the two networks as separate but coupled layers is what makes the representation useful. A graph that recorded only physical connections could say which terminals are linked, but not whether a journey between them survives once operators, timetables, and transfer windows are enforced. The functional network answers that second question, and because its structural layer is the simple-graph quotient of the supply multigraph in equation 3.11, the two are not separate objects but two readings of one system. A single disruption can therefore be scored on each, once by the connectivity it destroys and once by the feasible itineraries it removes.

Everything that follows runs on this one object. The single-mode analysis reads rail

and air robustness through it, the co-existence analysis asks whether the air layer offers usable alternatives to the rail layer within it, and the criticality and policy analysis ranks the links it exposes as most critical. Holding the topological representation beside the service representation is what later lets the difference between them be measured, the gap between what the network is structurally and what it can actually carry, which is where the rest of the thesis finds its answer.

Question 2 on unimodal service robustness - How robust are the European air and rail service networks under the random and targeted failure of their service links, namely the rail segments between consecutive stations and the non-stop flights that make up the supply layer? When the resulting service degradation is evaluated with service-level information, how do the results compare with topology-based robustness measures?

Both networks look robust under a topological or random assessment of failure and much less so under a service-level assessment, so evaluating the degradation with service-level information does not rescale the robustness estimate but moves what counts as critical. The contrast is sharpest on rail. The hub-and-spoke air network absorbs uncorrelated removals because alternative high-degree airports stand in for a lost one, whereas the long-distance rail network keeps its redundancy at the corridor rather than the station level and degrades more steeply as those corridors are thinned (Figure 5.2). Under targeted failure each mode concentrates its damage on a small set of edges, but service load and unweighted topology disagree about which set, and that disagreement is what answers the comparison the sub-question asks for. Integrated over the whole removal range the ordering holds, with the robustness-loss integral at $F = 0.934$ under the service-aware ordering against 0.891 under topology and 0.889 under random removal.

On the rail network the two rankings are effectively independent. The service-aware betweenness and the unweighted topology betweenness share no edges in their top fifty and only 4% in their top two hundred, and their full-edge-set rank correlation is $\rho = -0.03$ (tables 5.1 and 5.2). Topology elevates the sparse cross-border bridges that hold the continental graph in one piece, the Channel Tunnel approach, the Pyrenees gateway, and the Berlin to Poznań line, while service load elevates the dense German Rhine and Ruhr corridors and the Frankfurt to Mannheim main line through which most feasible itineraries actually run. Because the rail-only ranking is almost identical to the rail-dominated combined ranking ($\rho = 0.93$ and 0.94), this rail signature is also the one that governs the whole network.

On the air network the same divergence is present but alleviated by hub-and-spoke geometry. Almost every shortest path between two non-hub airports must cross a hub-to-hub link that also carries a large share of itineraries, so the two rankings agree more, overlapping on 17% of their top two hundred edges at a rank correlation of $\rho = 0.18$ (tables 5.3 and 5.4). Even here topology mislocates criticality, elevating degree-one cut-edges to peripheral airports while service load elevates the hub feeders into Heathrow, Stansted, Charles-de-Gaulle, and Schiphol and the Mediterranean island routes where the sea removes the rail option. The mechanism is identical on both modes, because unweighted betweenness counts geodesics through sparsely connected bridges and not the itineraries that travel the edge. What the two networks share is that the edges holding their graphs

together are not the edges carrying their journeys, strongly on rail and weakly on air, which is the first form of the gap the rest of the thesis measures.

Question 3 on rail-air co-existence - To what extent does rail-air co-existence enhance robustness at both the infrastructure and service layer, and which service links and corridors benefit most in terms of maintained connectivity under disruption?

Rail-air co-existence enhances robustness at the infrastructure layer and almost not at all at the service layer, and the distance between those two assessments is the answer. A structural assessment credits the combined network with cross-modal redundancy, and on a 124-city European rail and air network Ippolito and Cats (2024) find the air layer far more robust than rail, with a worst-case performance loss near 10% against 71%, and place the cities with the most redundant air links in a north-western European cluster spanning the Netherlands, Belgium, and western Germany. At the service layer that redundancy largely disappears. The combined retained-efficiency curve tracks the rail-only curve at every removal level (Figure 5.2), because rail-bearing itineraries dominate the feasible set and the pure-air alternatives are usually costlier per origin-destination pair and so add little retained efficiency, and the air layer's far greater stand-alone robustness does not transfer to the shared system.

The reason it does not transfer is mechanical, and the structural stress tests of Section 5.1 expose it. Service-aware failure removes the same dense rail corridors in the combined network as in the rail-only one, and each removed rail edge invalidates the pure-rail and the air-and-rail itineraries that cross it at once. Within the combined network at a 5% edge budget the air-involving itineraries do retain more than the rail-only ones, $R(x) = 0.344$ against $R(x) = 0.095$, yet 70.5% of them are already disconnected at that point, because the air leg reaches the network only through rail-side airport-feeder edges such as the Frankfurt Hbf to airport link and the Frankfurt to Mannheim to Stuttgart main line. Removing those feeders severs the air-and-rail itinerary at its rail end before the flight can act as a substitute. The combined network is amplified by the rail feeders rather than buffered by the air layer, and the 38.95% of itineraries that are rail-only set the floor on what survives.

The two macro-shocks of Section 5.3 test the same question from opposite directions and give the sharpest measurement of it. When the entire air mode closes, as in the Iceland 2010 ash scenario, the combined network barely moves and retains $R(x) = 0.875$ (Section 5.3.1), but the high figure is a composition effect rather than a cross-modal rescue. The air-to-rail substitution counter is flat zero, the 1,318 buckets that do record a cross-mode change are reroutes to a different open airport in the same urban area rather than a shift onto rail, and the rail-only itineraries that never traverse an airport place a floor under the retention on their own. When a flood collapses rail in its densest region, as in the July 2021 Bernd event, air does not measurably rescue it (Section 5.3.2), because roughly 81,000 intra-German buckets that keep a structurally feasible air substitute receive it outside the tolerance band and cannot use it. Co-existence therefore protects against air-side disruption far more than against rail-side disruption, and the narrow place where it does add a usable substitute is a small set of air-rail bridge airports, led by Frankfurt, where the ground link turns an air leg into a real alternative to a rail leg, so co-existence sharpens the existing criticality rather than redistributing it.

Even that narrow benefit is constrained by the rules under which a passenger may rebook, and the cooperation-tier experiment of Section 5.2.3 measures the constraint directly. Retained efficiency climbs across the tiers from $R(x) = 0.376$ under a single-operator ticket to $R(x) = 0.733$ under full rail interlining, and the final cross-modal step that adds the air mode on top of full rail interlining lifts it the rest of the way to the undisrupted baseline, a dividend of 0.267 retained-efficiency units (Figure 5.9). The service value of the second mode is therefore real and measurable, but it is realised only at the top tier of cooperation, and the present framework under Regulation 2021/782 (European Parliament and Council of the European Union, 2021) stands near the bottom of that curve with no cross-modal rebooking right at all.

Question 4 on planning and policy implications - Which service links and corridors are most critical for the robustness of the air network, the rail network, and the air-rail intermodal network respectively, and what targeted service or policy measures could mitigate the vulnerabilities identified?

The critical elements differ by mode, so no single protection list serves all three networks, and the largest mitigation is contractual rather than infrastructural. Each network concentrates its vulnerability on a structurally different class of element, on rail a dense set of national main corridors, on air a thin set of peripheral access edges, and on the combined network a handful of bridge airports. These three classes barely overlap, so protection directed at one network confers little benefit on the others.

On the rail network criticality concentrates on the German Rhine-Ruhr to Frankfurt-Mannheim corridors, which carry most of the feasible itineraries, with 88% of the top two hundred service-aware edges German, alongside a separate topological tier of cross-border gateway links through which national systems reach the continental core. The single-operator withdrawal of Deutsche Bahn alone leaves $R(x) = 0.390$ with 54.4% of OD-departure buckets disconnected (Section 5.2.1). This corridor is already known outside this thesis as Europe's principal rail bottleneck, the constrained mixed passenger-freight axis that Otsuka et al. (2017) trace through Düsseldorf, Cologne, Frankfurt, and Mannheim, the region where Witlox et al. (2022) find high-speed services losing operating speed, and its cities appear among the routes on which rail is already faster than air in the door-to-door mapping of Brons et al. (2023). The contribution here is to show that the corridor which is congested in normal operation is also the point at which the network is least able to absorb a disruption, so its criticality is operational and structural at once. Air inverts this signature. Its mega-hubs carry little criticality because comparable hubs absorb their removal, and the exposure falls instead on the peripheral airports that hold the only non-stop link to their region, such as Brno, Poprad-Tatry, and Rijeka with a single scheduled connection each in the March timetable, and on the hub-feeder funnels into Heathrow, Stansted, Charles-de-Gaulle, and Schiphol. The combined network concentrates criticality on a third class again, the few air-rail bridge airports led by Frankfurt where a ground link turns an air leg into a usable rail substitute.

Three targeted measures follow from this geography. First, the protection of physical assets must be set per mode and per criterion, since the service-critical German corridor and the topology-critical cross-border gateways are different elements, so a robustness programme needs a dedicated concentration-protection funding line scoped to the German corridor alongside the cross-border connectivity funding that the next

trans-European network cycle already carries (European Parliament and Council of the European Union, 2024b). Second, because unweighted topology betweenness points at the gateway links and misses the corridors that actually carry the itineraries, service-aware centrality is the more reliable planning basis, the same mismatch that Pien et al. (2015) measured between betweenness and operational delay on the European air network. Third, because the macro-shocks show co-existence protecting against air-side disruption far more than against rail-side disruption, the rail-side response to an event on the German corridors is surge capacity for the stranded passengers rather than an expectation of cross-modal rescue, and climate-adaptation budgets should follow the country-specific exposure rather than a uniform standard.

The larger mitigation, however, is contractual, as the answer to the main research question above quantifies. How far a passenger may rebook across operators and modes matters more than the infrastructural budget of the worst disruption. The dominant steps are the move to alliance-tier rail interlining and the final cross-modal step that adds the air mode, and the present framework under Regulation 2021/782 (European Parliament and Council of the European Union, 2021) captures almost none of either, providing limited cross-operator rail rights and no cross-modal equivalent. Strengthening it into a binding cross-operator and cross-modal rebooking default, in line with the cross-border passenger-rail action plan of European Commission (2021), would deliver more robustness on the same timeline than any single infrastructure project, because it converts the alternatives that already exist in the network into journeys a passenger can use. The most effective protection for the German corridor is therefore not more track but a rule that lets a disruption on it be rebooked across every operator and mode that crosses it.

This section sets out the limitations of the study and the directions in which future work could extend it.

6.2. Limitations

Each limitation below is stated with what it affects, what it leaves intact, and the data or method by which a later study could address it.

The analysis evaluates the availability and generalised cost of feasible itineraries, not the number of travellers carried or the welfare they lose. The functional network of Section 3.1.4 carries scheduled services rather than passenger flows, so it reports how the supply of usable itineraries degrades, not how many travellers are affected (Section 3.1.1). This bounds the policy interpretation, because a corridor that carries many alternatives is not automatically the one that carries the most travellers. It leaves the central comparison intact, because the divergence between the topological and service-aware perspectives of Section 5.1 is a property of the supply graph, independent of any demand weighting. This omission could be addressed with data beyond the present scope, since automated fare-collection or ticketing records would supply the origin-destination flows needed to weight the enumerated itineraries by travellers, in the demand-sensitive tradition of Cats and Jenelius (2014).

The measure that ranks those itineraries is a comparative impedance, not a behavioural model. The generalised travel cost of equation 3.21 carries no alternative-specific constants and no mode-specific perception terms, so the cheapest enumerated itinerary is not claimed to be the one a given traveller would choose. This bounds the behavioural inter-

pretation of any single reroute, but not the ranking of corridors, because the cost is applied consistently to every alternative and the sensitivity analysis of Section 5.4.5 confirms that the orderings hold across the plausible range of weights. A retained alternative is likewise counted as available whenever it stays feasible under the transfer and tolerance rules of Section 3.2.3, whether or not the surviving services hold the seats to carry the displaced travellers, so the retained efficiency of equation 3.25 is an upper bound that a capacity constraint could only lower. A later study could calibrate the cost to revealed or stated preferences and draw seat and slot capacities from booking data, converting both bounds into operational estimates and sizing the reserve capacity that Cats and Jenelius (2015) identify as the response to unforeseen disruption.

Two further limitations follow from the data. The network is built from a fixed schedule and a disruption is applied as an instantaneous removal through the disruption rule of Section 3.3.3, so the framework measures robustness at the moment of disruption, not the recovery that follows, and the air schedule of Chapter 4 is a single representative period rather than a continuous series. The reported robustness figures are correspondingly steady-state averages over the enumeration window rather than a time-resolved series. The appendix stratifies the rail pool by departure hour and finds that the hour of travel lowers the retained-efficiency curve by up to eight points at heavy removal (Figure 6.2). Time of day nonetheless explains only about one per cent of the variation, so the steady-state figure remains the correct central estimate. A comparison across calendar time is not possible, because the national rail feeds cover the 2024 to 2026 window unevenly (Figure 3.6), so any slice by month, season, or year would track feed coverage rather than real change in the network. Additionally, coverage also differs across countries, several appearing on the air layer only or not at all, and the rail graph of equation 3.5 is built from ordered stop sequences rather than track geometry, so parallel tracks are not represented separately and the reported times are station-to-station rather than door-to-door. Neither limitation changes which corridors and operators are critical, because every ranking is evaluated at the same disruption instant, and the German network that Section 6.1.1 identifies as most critical is also the most completely covered part of the system. Real-time operations data would allow a successor to follow the recovery this snapshot omits, in the rescheduling tradition reviewed by Cacchiani et al. (2014), and richer feeds with track geometry would widen the coverage and restore the door-to-door legs.

6.3. Implications for future study

The complex-network analysis of transport systems is valued for the generality of its topological abstraction, since a single graph representation makes networks of different scale and geography comparable (Barabási, 2016; von Ferber et al., 2009). That abstraction can miss what an operator or a traffic engineer needs, because a purely topological account carries no record of how service is organised on the graph. This divergence between structure and operation is what Bešinović (2020) review for rail, and what Cats and Jenelius (2014) address by weighting each link with the flow it carries.

The methodological contribution of this thesis is a framework built to answer that concern, and it is this framework, more than any single result, that a successor inherits. It holds the same system at once as a structural network of connections and as a functional network, the second overlaying operator, timetable, and generalised travel cost on every

scheduled service. Topological and service-aware robustness are then measured at the same scale, so their divergence is a result rather than an artefact of two datasets. On that basis a service-aware criticality measure ranks each link by the journeys that depend on it, one disruption rule spans structural, operator, and macro-scenario failures, and a four-state classifier records how service degrades. The framework transfers unchanged to any continental system whose scheduled services can be enumerated.

The divergence this framework exposes, that a topological account overstates the service a traveller can actually obtain, is not confined to robustness. Accessibility is the clearest case, and the introduction already marks it as an extension of the same basis. Network-science accounts of how easily a place can be reached derive it from the same topological connectivity and inherit the same gap (Luo et al., 2019). Recomputing accessibility on the service layer, with generalised travel cost in place of graph distance, would relocate which places are genuinely reachable, as it here relocated which links are genuinely critical, and could revise the standing of the peripheral regions a topological measure rates as well connected.

The same shift bears on how the modes are compared. Where accounts of air-rail substitution weigh the two by distance (Bruno et al., 2025; Dobruszkes, 2011), the enumerated itineraries identify the corridors on which a passenger can in fact choose between them, and the finding that cross-modal substitution rarely yields a usable alternative sets a service-level ceiling on how much traffic a mode-shift policy could move. A natural next step brings in the road modes left outside this study, the private car and long-distance coach, to test whether they supply the fallback that rail and air do not. Continental robustness studies that measure criticality from the graph alone (Cardillo et al., 2013; Ippolito & Cats, 2024) invite the same re-examination, since the service layer may relocate their critical elements as it did here.

A further direction turns robustness from a property to be measured into an objective to be designed. The retained efficiency of this study can serve as the quantity to optimise, seeking the timetable, service pattern, or cooperation arrangement that minimises the loss a given disruption inflicts. The cooperation set C, held fixed here, then becomes a design variable rather than a context, so that the coalitions forming among the 51 rail operators and 31 carriers, and the passenger-rights rules that govern them, are the coalition-formation and policy-design questions that Cats (2025) name as open, which the game-theoretic setting of Adler et al. (2010) equips and the fragmented open-access market of Beria et al. (2023) motivates. In each of these directions a successor can start from the framework of this thesis rather than beside it.

Appendix

Inter-terminal transfer

Table 6.2: The sixty most-traversed inter-terminal connecting legs in the enumerated path pool, with the urban area, the two terminals, the mode-pair, the connecting service and its time at the midday band, and the traversal count. Time t in minutes, traversals in millions of path-alternative legs.

urban area	terminal 1	terminal 2	pair	connecting service	t	uses
Amsterdam	Schiphol Airport	Amsterdam Schiphol (AMS)	rail-air	walk	6	24.7
Düsseldorf	Düsseldorf Hbf	Düsseldorf (DUS)	rail-air	heavy rail, cable car	24	19.7
Frankfurt am Main	Frankfurt (M) Hbf	Frankfurt (FRA)	rail-air	conventional rail	42	18.8
Düsseldorf	Duisburg Hbf	Düsseldorf (DUS)	rail-air	commuter rail, heavy rail	35	12.8
Amsterdam	Rotterdam Central	Amsterdam Schiphol (AMS)	rail-air	heavy rail	34	12.1
Frankfurt am Main	Flughafen Fernbf	Frankfurt (FRA)	rail-air	bus	37	10.4
München	München Hbf	München (MUC)	rail-air	bus	54	9.3
Amsterdam	Utrecht Centraal	Amsterdam Schiphol (AMS)	rail-air	heavy rail	42	8.2
Berlin	Berlin Hbf	Berlin Brandenburg (BER)	rail-air	commuter rail	49	7.8
Amsterdam	Den Haag Laan v NOI	Amsterdam Schiphol (AMS)	rail-air	heavy rail	33	7.6
Hamburg	Hamburg Hbf	Hamburg (HAM)	rail-air	commuter rail	45	7.3
Amsterdam	Amsterdam Centraal	Amsterdam Schiphol (AMS)	rail-air	heavy rail	26	7.2
Paris	Charles-de-Gaulle Airport	Paris - Charles de Gaulle (CDG)	rail-air	bus	20	7.1
Frankfurt am Main	Mainz Hbf	Frankfurt (FRA)	rail-air	commuter rail	55	6.9
Berlin	Berlin Ostbahnhof	Berlin Brandenburg (BER)	rail-air	commuter rail	37	6.6
Amsterdam	The Hague Central	Amsterdam Schiphol (AMS)	rail-air	heavy rail	52	5.4
Zurich	Olten	Zürich (ZRH)	rail-air	long distance train	67	5.1
Düsseldorf	Düsseldorf Flughafen	Düsseldorf (DUS)	rail-air	bus	15	5.0
Frankfurt am Main	Mannheim Hbf	Frankfurt (FRA)	rail-air	conventional rail	62	4.9
Zurich	Zürich Airport	Zürich (ZRH)	rail-air	walk	13	4.9
Luton	Luton	London-Luton (LTN)	rail-air	bus	14	4.8
London	London Paddington	London-Heathrow (LHR)	rail-air	heavy rail	40	4.7

continued on next page

Table 6.2 (continued)

urban area	terminal 1	terminal 2	pair	connecting service	t	uses
München	München Pasing	München (MUC)	rail-air	heavy rail	63	4.7
Hamburg	Hamburg Dammtor	Hamburg (HAM)	rail-air	bus, metro	40	4.5
Amsterdam	Amsterdam Sloterdijk	Amsterdam Schiphol (AMS)	rail-air	heavy rail	19	3.8
Düsseldorf	Düsseldorf Flughafen Terminal	Düsseldorf (DUS)	rail-air	walk	3	3.8
Düsseldorf	Düsseldorf Wehrhahn	Düsseldorf (DUS)	rail-air	commuter rail	12	3.7
Nürnberg	Nürnberg Hbf	Nürnberg (NUE)	rail-air	metro	24	3.7
Mulhouse	Olten	Mulhouse-Bale (BSL)	rail-air	bus, long distance train	56	3.6
Frankfurt am Main	Langen (Hessen)	Frankfurt (FRA)	rail-air	commuter rail, bus	50	3.6
Stuttgart	Stuttgart Hbf	Stuttgart (STR)	rail-air	metro, bus	71	3.6
Bremen	Bremen Hbf	Bremen (BRE)	rail-air	tram	23	3.5
Amsterdam	Den Haag HS	Amsterdam Schiphol (AMS)	rail-air	heavy rail	36	3.1
Paris	Paris Gare de Lyon	Paris - Charles de Gaulle (CDG)	rail-air	bus, commuter rail	61	3.0
Frankfurt am Main	Darmstadt Hbf	Frankfurt (FRA)	rail-air	bus	53	3.0
Düsseldorf	Cologne Hbf	Düsseldorf (DUS)	rail-air	heavy rail, conventional rail	46	2.9
London	Reading	London-Heathrow (LHR)	rail-air	heavy rail	72	2.9
Hannover	Hannover Hbf	Hannover (HAJ)	rail-air	commuter rail	32	2.8
Lyon	Lyon Part Dieu	Lyon Saint-Exupery (LYS)	rail-air	tram	37	2.8
Düsseldorf	Cologne Messe/Deutz	Düsseldorf (DUS)	rail-air	heavy rail, conventional rail	53	2.7
Berlin	Berlin Gesundbrunnen	Berlin Brandenburg (BER)	rail-air	commuter rail	61	2.7
Paris	Paris Est	Paris - Charles de Gaulle (CDG)	rail-air	bus, commuter rail	56	2.7
Stockholm	Stockholm Central	Stockholm - Arlanda (ARN)	rail-air	conventional rail	36	2.6
Köln	Troisdorf	Köln/Bonn (CGN)	rail-air	commuter rail	46	2.6
Kraków	Kraków	Krakow Balice (KRK)	rail-air	heavy rail	21	2.6
Amsterdam	Amsterdam Zuid	Amsterdam Schiphol (AMS)	rail-air	heavy rail	15	2.5
Düsseldorf	Mönchengladbach Hbf	Düsseldorf (DUS)	rail-air	commuter rail, heavy rail	49	2.5
Frankfurt am Main	Mainz Römisches Theater	Frankfurt (FRA)	rail-air	commuter rail	53	2.5
London	London Euston	London-Heathrow (LHR)	rail-air	heavy rail, metro	56	2.5
Dresden	Dresden Neustadt	Dresden (DRS)	rail-air	heavy rail	14	2.4
Düsseldorf	D-Unterrath S	Düsseldorf (DUS)	rail-air	heavy rail	4	2.4
Düsseldorf	Troisdorf	Düsseldorf (DUS)	rail-air	commuter rail, heavy rail, cable car	79	2.3
Zurich	Lenzburg	Zürich (ZRH)	rail-air	heavy rail	56	2.3

continued on next page

Table 6.2 (continued)

urban area	terminal 1	terminal 2	pair	connecting service	t	uses
Amsterdam	Breukelen	Amsterdam Schiphol (AMS)	rail-air	heavy rail	42	2.3
Paris	Paris Montparnasse	Paris - Charles de Gaulle (CDG)	rail-air	bus, commuter rail, metro	75	2.3
Paris	Paris Nord	Paris - Charles de Gaulle (CDG)	rail-air	bus, commuter rail	53	2.3
Köln	Cologne Airport	Köln/Bonn (CGN)	rail-air	walk	43	2.3
Bristol	Bristol Temple Meads	Bristol (BRS)	rail-air	bus	40	2.2
Bruxelles / Brussel	Brussels North	Brussels National (BRU)	rail-air	heavy rail	18	2.2
Edinburgh	Haymarket	Edinburgh (EDI)	rail-air	bus	26	2.1

Single-operator withdrawal catalogue

Table 6.3 reports the full single-operator withdrawal catalogue that accompanies the operator-removal experiment of Section 5.2, of which only the ten most damaging operators are tabulated in the main text.

Table 6.3: Full $k = 1$ single-operator withdrawal catalogue on the multimodal pool, all 82 operators (51 rail plus 31 air), sorted by retained efficiency $R(x)$ (most damaging first). The edge count is the operator supply-layer footprint on the 6,760-edge intermodal pool and x its edge-removal-equivalent intensity. Values are taken from the no-bus $k = 1$ criticality arrays.

Operator	Mode	Edges	x	$R(x)$	f_{disc} (%)	$\Delta\bar{D}$ (min)
DB	rail	1499	0.222	0.390	54.4	60.8
NS	rail	200	0.030	0.890	12.9	18.8
SNCF	rail	407	0.060	0.919	11.5	14.8
NMBS/SNCB	rail	143	0.021	0.919	8.0	10.3
DFTO	rail	291	0.043	0.936	6.6	9.3
ÖBB	rail	377	0.056	0.949	4.4	4.7
FR	air	1364	0.202	0.950	9.7	19.1
SBB	rail	269	0.040	0.956	7.5	12.3
Trenitalia	rail	260	0.038	0.961	5.2	7.5
U2	air	479	0.071	0.971	5.4	11.7
GTR	rail	127	0.019	0.974	3.6	5.6
KL	air	80	0.012	0.976	5.1	10.5
EW	air	164	0.024	0.979	4.5	9.5
LH	air	116	0.017	0.979	5.2	10.5
CrossCountry	rail	89	0.013	0.979	1.4	1.9
FirstGroup	rail	107	0.016	0.982	2.2	3.3
Eurostar	rail	24	0.004	0.985	1.0	1.6
LX	air	81	0.012	0.985	3.5	6.9
PKP	rail	114	0.017	0.986	2.4	4.0
BA	air	102	0.015	0.987	2.7	6.0
Avanti West Coast	rail	72	0.011	0.988	1.1	1.8
AF	air	77	0.011	0.990	2.1	4.4
Renfe	rail	84	0.012	0.993	1.6	2.8
West Midlands Trains	rail	23	0.003	0.994	0.6	0.8
MÁV	rail	73	0.011	0.994	1.5	2.6
PolRegio	rail	97	0.014	0.994	0.9	1.4
ScotRail	rail	36	0.005	0.995	0.9	1.4
Trenord	rail	81	0.012	0.995	1.1	1.9

Table 6.3 continued from previous page

Operator	Mode	Edges	x	$R(x)$	f_{disc} (%)	$\Delta\bar{D}$ (min)
SOB	rail	36	0.005	0.996	0.8	1.3
W6	air	154	0.023	0.996	1.0	2.0
ZSSK	rail	58	0.009	0.996	0.8	1.1
SJ	rail	24	0.004	0.996	0.9	1.6
LO	air	56	0.008	0.997	0.7	1.4
SN	air	39	0.006	0.997	0.7	1.5
HV	air	52	0.008	0.997	0.6	1.2
IB	air	74	0.011	0.997	0.6	1.3
SK	air	64	0.009	0.997	0.6	1.2
V7	air	102	0.015	0.997	0.6	1.3
AZ	air	48	0.007	0.997	0.5	1.2
OS	air	61	0.009	0.997	0.5	1.0
DSB	rail	22	0.003	0.998	0.3	0.4
AY	air	44	0.007	0.998	0.3	0.7
LG	air	34	0.005	0.998	0.3	0.7
CFL	rail	20	0.003	0.998	0.2	0.2
BT	air	63	0.009	0.998	0.3	0.5
DY	air	59	0.009	0.998	0.3	0.6
TO	air	47	0.007	0.999	0.3	0.7
WESTbahn	rail	10	0.001	0.999	0.1	0.2
LS	air	109	0.016	0.999	0.3	0.6
A3	air	58	0.009	0.999	0.2	0.4
Grand Central	rail	22	0.003	0.999	0.1	0.2
CP	rail	13	0.002	0.999	0.2	0.4
TP	air	42	0.006	0.999	0.1	0.3
QS	air	37	0.005	0.999	0.1	0.3
OU	air	18	0.003	0.999	0.1	0.2
RO	air	15	0.002	0.999	0.1	0.2
CFR	rail	47	0.007	0.999	0.1	0.2
Iarnród Éireann	rail	4	0.001	0.999	0.1	0.3
VR	rail	9	0.001	0.999	0.1	0.2
Snälltåget	rail	17	0.003	1.000	0.1	0.1
Vy	rail	12	0.002	1.000	0.1	0.2
European Sleeper	rail	11	0.002	1.000	0.0	0.1
FB	air	16	0.002	1.000	0.1	0.1
Go-Ahead	rail	22	0.003	1.000	0.0	0.0
KM	air	16	0.002	1.000	0.1	0.1
Ouigo España	rail	8	0.001	1.000	0.0	0.1
VR Snabbtåg	rail	4	0.001	1.000	0.0	0.1
CD	rail	20	0.003	1.000	0.0	0.0
LTG Link	rail	2	0.000	1.000	0.0	0.0
SJ Nord	rail	4	0.001	1.000	0.0	0.0
Go-Ahead Norge	rail	2	0.000	1.000	0.0	0.0
HZZP	rail	11	0.002	1.000	0.0	0.0
Transferoviar	rail	6	0.001	1.000	0.0	0.0
Regio Călători	rail	7	0.001	1.000	0.0	0.0
Tågab	rail	6	0.001	1.000	0.0	0.0
Astra Trans Carpatic	rail	3	0.000	1.000	0.0	0.0
Softrans	rail	2	0.000	1.000	0.0	0.0
Elron	rail	1	0.000	1.000	0.0	0.0
Caledonian Sleeper	rail	38	0.006	1.000	0.0	0.0
EMR	rail	34	0.005	1.000	0.0	0.0
Hellenic Train	rail	0	0.000	1.000	0.0	0.0
FI	air	0	0.000	1.000	0.0	0.0

Edge-removal sweeps

Tables 6.4 and 6.5 report the retained efficiency $R(x)$ at each edge-removal fraction x behind the curves of Section 5.1, read from the same no-bus arrays as the chapter's headline values.

Table 6.4: Retained efficiency $R(x)$ on the intermodal pool at each removal fraction x, under the service-aware and topology betweenness orderings and the 30-seed random benchmark (median). Source of figures 5.2 and 5.4.

x	service-aware	topology	random
0.00	1.000	1.000	1.000
0.05	0.168	0.557	0.577
0.10	0.100	0.292	0.353
0.15	0.067	0.183	0.220
0.20	0.051	0.122	0.150
0.25	0.042	0.086	0.103
0.30	0.037	0.057	0.076
0.35	0.035	0.045	0.055
0.40	0.033	0.040	0.042
0.45	0.033	0.037	0.032
0.50	0.032	0.034	0.025
0.55	0.032	0.032	0.020
0.60	0.031	0.030	0.015
0.65	0.029	0.028	0.012
0.70	0.027	0.026	0.009
0.75	0.025	0.024	0.007
0.80	0.023	0.022	0.005
0.85	0.020	0.020	0.004
0.90	0.018	0.018	0.002
0.95	0.017	0.016	0.001
1.00	0.014	0.014	0.000

Table 6.5: Retained efficiency $R(x)$ under the 30-seed random benchmark (median) on the pure-rail, pure-air, and intermodal pools. The air network degrades far more slowly than the rail network, the intrinsic mode-robustness contrast of Section 5.1.2.

x	rail	air	intermodal
0.00	1.000	1.000	1.000
0.05	0.513	0.931	0.577
0.10	0.272	0.864	0.353
0.15	0.154	0.802	0.220
0.20	0.096	0.735	0.150
0.25	0.062	0.671	0.103
0.30	0.039	0.608	0.076
0.35	0.027	0.546	0.055
0.40	0.019	0.490	0.042
0.45	0.014	0.428	0.032
0.50	0.010	0.375	0.025
0.55	0.007	0.319	0.020
0.60	0.005	0.271	0.015
0.65	0.004	0.223	0.012
0.70	0.003	0.180	0.009
0.75	0.002	0.140	0.007
0.80	0.001	0.103	0.005
0.85	0.001	0.071	0.004
0.90	0.000	0.044	0.002
0.95	0.000	0.019	0.001
1.00	0.000	0.000	0.000

Cooperation-tier results

Table 6.6 reports the retained efficiency of the cooperation-tier restriction of Section 5.2.3, and Table 6.7 the seven documented bilateral pairs that make up tier B2.

Table 6.6: Retained efficiency $R(x)$ across the five cooperation tiers of Section 5.2.3 on the intermodal pool, with the cooperation-set size $|C|$. The span from a single-operator ticket to full modal interline is 0.624 retained-efficiency units. Companion of Figure 5.9.

Tier	Cooperation arrangement	$\lvert C\rvert$	$R(x)$
B1-strict	single-operator ticket (Deutsche Bahn)	1	0.376
B1-split	one operator change	2	0.426
B2	best documented bilateral pair	2	0.426
B3	Railteam alliance	9	0.631
B4	full rail interline	52	0.733
B5	full modal interline	83	1.000

Table 6.7: The seven documented bilateral cooperation pairs of tier B2 and their retained efficiency $R(x)$ on the intermodal pool. The four Deutsche Bahn-anchored pairs reach $R(x)$ near 0.4, while the three pairs without Deutsche Bahn fall close to zero because each covers only a single international relation.

Pair	Relation	$R(x)$
DB and ÖBB	Nightjet	0.426
DB and NMBS	Aachen-Liège	0.414
DB and SNCF	Paris-Frankfurt-Stuttgart	0.410
DB and SBB	Munich-Zürich	0.391
SNCF and Trenitalia	Paris-Milan	0.028
SNCF and Renfe	Madrid-Lyon	0.022
Eurostar	post-Thalys restructuring	0.001

Macro-shock scenarios

Table 6.8 reports the retained efficiency and disconnection share of the Type III macro-shock scenarios of Section 5.3. The ash values are read from the no-bus ash arrays; the Bernd flood values are the intermodal-pool figures of Section 5.3.2, whose per-country breakdown is in Table 5.6.

Table 6.8: Retained efficiency $R(x)$ and disconnected share f_{disc} for the Type III macro-shock scenarios. The ash hazard is an air-only footprint and the Bernd flood a rail-side footprint.

Scenario	$R(x)$	f_{disc} (%)
Iceland 2010 ash, peak day (153 airports)	0.875	20.6
Iceland 2010 ash, five-zone union (178 airports)	0.853	24.5
Bernd 2021 flood, Europe-wide scope	0.679	–
Bernd 2021 flood, intra-German scope	0.510	47.8

Betweenness rankings

Tables 6.9 and 6.10 list the top edges of the rail and air networks under the two betweenness orderings of Section 5.1.1, the service-aware betweenness $\hat{b}$ and the unweighted topology betweenness b^{topo}. The two orderings select almost disjoint sets of edges, the service-aware ordering elevating the dense national corridors that carry the itineraries and the topology ordering the sparse cross-border links that hold the graph together. The intermodal service-aware ranking is dominated by and near-identical to the rail ranking ($\rho = 0.94$, Section 5.1.1) and is therefore represented by the rail table.

Table 6.9: The eight highest rail edges under the service-aware betweenness $\hat{b}$ and the topology betweenness b^{topo}, on the pure-rail network. The two rankings share no edge.

	service-aware edge	$\hat{b}$ ($\times 10^3$)	topology edge	b^{topo}
1	Düsseldorf Hbf – Duisburg Hbf	132	Brussel-Zuid – St. Pancras	0.164
2	Essen Hbf – Duisburg Hbf	126	Cologne Hbf – Brussel-Zuid	0.080
3	Frankfurt Hbf – Flughafen Fernbf	109	St. Pancras – Finsbury Park	0.070
4	Ulm Hbf – Augsburg Hbf	103	Barcelona Sants – Perpignan	0.065
5	Stuttgart Hbf – Ulm Hbf	98	Lyon Part Dieu – Perpignan	0.064
6	Hamburg Hbf – Hamburg Dammtor	96	St. Pancras – Selby	0.063
7	München Pasing – München Hbf	94	St. Pancras – London Blackfriars	0.059
8	Hamburg Hbf – Hamburg Harburg	90	Brussel-Zuid – Amsterdam Centraal	0.052

Table 6.10: The eight highest air edges under the service-aware betweenness $\hat{b}$ and the topology betweenness b^{topo}, on the pure-air network. The topology ordering ties the eleven degree-one cut-edges at $b^{\text{topo}} = 7.4 \times 10^{-3}$, each the only non-stop link to a peripheral airport.

	service-aware edge	$\hat{b}$	topology edge	b^{topo}
1	ABZ – LHR	883	ERF – PMI	0.0074
2	BHD – LHR	798	BRQ – STN	0.0074
3	GLA – LHR	758	HEL – OUL	0.0074
4	LHR – NCL	745	MAD – PNA	0.0074
5	CDG – MPL	693	ORB – STN	0.0074
6	AMS – BLL	680	HEL – JYV	0.0074
7	DUB – STN	677	STN – TAT	0.0074
8	AAL – AMS	667	OPO – XCR	0.0074

Temporal coverage of the input feeds and its effect on robustness

The robustness results reported in this thesis are steady-state averages over the full enumeration window. Two input feeds bound how far a temporal view of those results can be trusted. The air schedule covers March 2025 only and is carried across every rail-operating day, so a seasonal view of the intermodal pool cannot separate the air-coverage limit from real structure. The rail date feed, examined below, is itself a publish-and-thin forward snapshot, so a month-of-year view of rail is an operator-data-coverage matter rather than a seasonal robustness signal. Departure hour is the one feed-internal axis that is clean for rail, and stratifying by it bounds how much an hour-resolved curve would depart from the reported aggregate. The rail pool examined holds 123,201,824 alternatives over 1,649 stations.

The rail date feed does not cover the three years uniformly. Through 2024 it carries about 22 distinct services per month, which is a few long-validity services rather than a represented network. It then ramps to a dense plateau of roughly thirty to forty-seven thousand services from November 2025 to May 2026, after which it thins toward seven thousand services by the end of 2026 as fewer trips carry validity that far forward. The enumerated pool tracks this gradient one to one, from a few hundred station pairs in 2024 to more than three hundred thousand at the December 2025 peak. A robustness curve sliced by month would therefore rise and fall with feed coverage and not with any property of European rail operation. This is the rail counterpart of the air-coverage limit

and is the reason the analysis reports month-blind aggregates.

Departure hour, by contrast, varies within the dense window and reflects real service. The set of station pairs served swings by a factor of 15.9 across the day, from 22,685 pairs at 23:00 to 361,028 pairs at 06:00. Among served pairs, the share carried by two alternatives or fewer, which is the share that fails first under any removal of two edges or two routes, ranges from 0.122 in the early evening to 0.247 late at night.

To convert this gradient into a robustness number, each station pair with k alternatives is given a survival probability under independent random removal of a fraction x of its alternatives,

$$\tilde{R}(x) \;=\; \mathbb{E}_{\mathrm{OD}}\big[\,1-x^{k}\,\big], \tag{6.1}$$

averaged over station pairs, where x here is the fraction of an origin-destination pair's own alternatives removed rather than the edge-removal fraction used in the main analysis. This is a redundancy-only proxy that ignores edge sharing and uses random rather than targeted removal, so it is a lower bound on the movement a targeted edge-removal sweep would produce.

Table 6.11 reports the outcome. The reported aggregate pools all departure times and so carries the largest effective redundancy, and it sits above every single hour. The spread between the best and worst hour grows from 0.008 at $x = 0.10$ to 0.055 at $x = 0.50$, and the worst hour, 23:00, sits below the reported aggregate by 0.011 at $x = 0.10$ and by 0.085 at $x = 0.50$. Figure 6.2 shows the envelope. A passenger constrained to a single departure hour, in the late-night or mid-morning trough, meets robustness up to about eight points below the all-hours figure, while the all-hours figure stays the correct central estimate for a passenger free to choose a departure time.

Two statistics qualify this shift under the data limit. With one hundred thousand to one million station pairs per hour, the standard error of $\tilde{R}(x)$ is of order 10^{-4}, so the per-hour confidence intervals are about 10^{-3} wide and do not overlap, and a two-sample z between the worst hour and the pooled aggregate reaches 50.8 at $x = 0.10$ and 75.4 at $x = 0.50$. The effect is statistically real, so the data limit is not a matter of statistical power. The effect size is nonetheless small. The fraction of variation that departure hour explains is the one-way effect size η^{2}, defined for the station-pair survival $s_i(x) = 1 - x^{k_i}$ of origin-destination pair i with k_i alternatives, grouped by departure hour h, as

$$\eta^{2}(x) \;=\; \frac{\sum_{h} n_h \left(\bar{s}_h(x) - \bar{s}(x)\right)^{2}}{\sum_{i} \left(s_i(x) - \bar{s}(x)\right)^{2}}, \tag{6.2}$$

with $\bar{s}_h(x)$ the mean survival within hour h, $\bar{s}(x)$ the grand mean over all pairs, and n_h the number of pairs departing in hour h, so the numerator is the between-hour variation and the denominator the total. This $\eta^{2}(x)$ rises only from 0.005 at $x = 0.10$ to 0.011 at $x = 0.50$, so time of day explains about one percent of the variation in station-pair robustness and heterogeneity in alternative count across station pairs explains the remaining ninety-nine percent. Path redundancy across station pairs, not the time of day, is the dominant driver.

The year-month axis fails for a different statistical reason, namely identification rather than power. Across months the pool coverage and the feed service count move together with a log-scale Pearson correlation of 0.984, and the coefficient of variation of pool coverage is 2.84 across all months against 1.10 inside the dense plateau. A month-level ro-

Table 6.11: Departure-hour shift in the redundancy proxy of Equation (6.1). The aggregate is the all-hours pooled value. The worst and best columns are over the 24 clock hours. The gap is the aggregate minus the worst hour.

x	aggregate $\tilde{R}(x)$	worst hour	best hour	aggregate − worst
0.05	0.9987	0.9934	0.9971	0.0053
0.10	0.9973	0.9862	0.9939	0.0111
0.20	0.9938	0.9694	0.9861	0.0244
0.30	0.9892	0.9487	0.9760	0.0405
0.50	0.9749	0.8902	0.9451	0.0847

bustness figure would inherit feed coverage almost entirely, so a month-of-year contrast would measure data coverage and not a seasonal property of the network.

Three consequences follow. The critical-edge ranking is hour-invariant in practice, because the leading corridors are trunk links that carry departures in every hour, so the identity of the most critical edges does not change under stratification. The absolute level of the robustness curve is hour-sensitive, and an honest band on the reported rail figure is a downward shift of up to about eight points at heavy removal for a single departure hour. The intermodal figures carry both the air-coverage limit and the rail feed-coverage limit, so a seasonal intermodal view is doubly unavailable, and the steady-state framing is the defensible one for both modes under a shared assumption set.

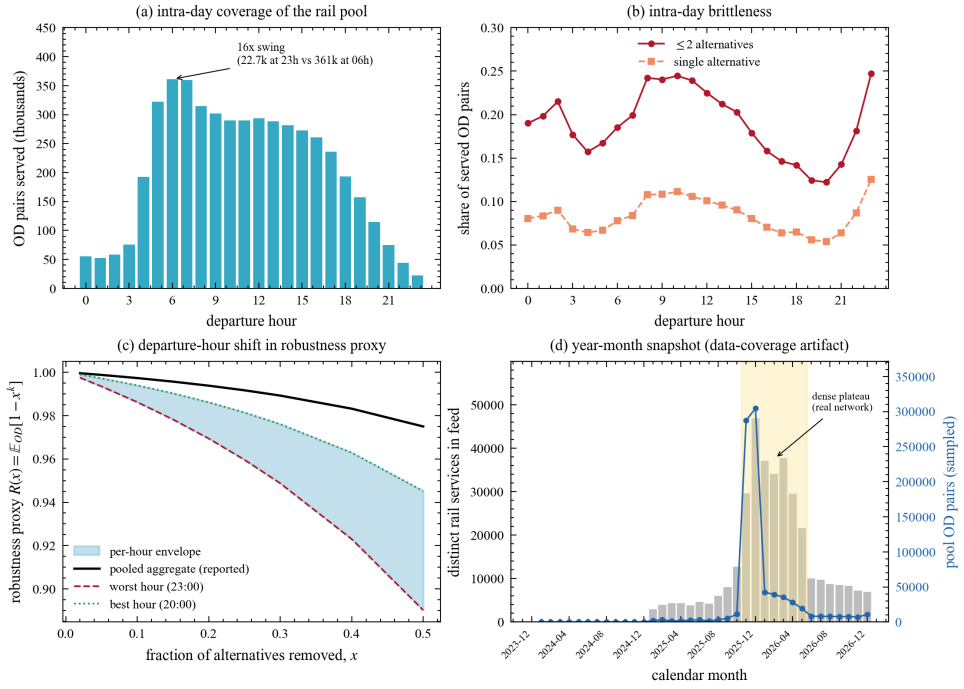

Figure 6.2: Temporal stratification of the rail itinerary pool. Panel (a) shows the count of station pairs served by clock hour. Panel (b) shows the share of served station pairs carried by two alternatives or fewer and by a single alternative. Panel (c) shows the redundancy proxy of Equation (6.1), the reported all-hours aggregate against the per-hour envelope and the worst and best hours. Panel (d) is the year-month snapshot, feed service counts with pool station-pair coverage, shaded over the dense plateau, shown as a data-coverage matter rather than a seasonal signal.

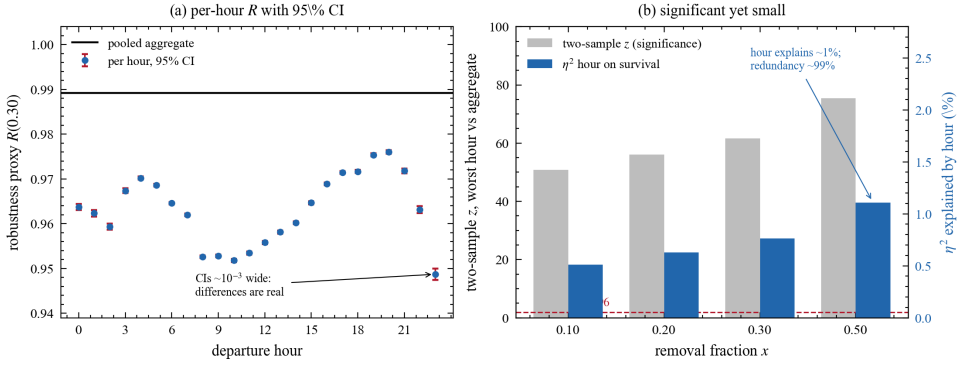

Figure 6.3: Significance against effect size for the departure-hour stratification. Panel (a) shows the per-hour redundancy proxy $\tilde{R}(x)$ at $x = 0.30$ with 95 percent confidence intervals about 10^{-3} wide, which do not overlap across hours, against the pooled aggregate. Panel (b) places the two-sample z between the worst hour and the aggregate, which reaches about 75 at heavy removal, beside the variance fraction η^2 of equation (6.2), which stays near one percent. The departure-hour effect is statistically unambiguous and substantively small.

Bibliography

Adler, N., Pels, E., & Nash, C. (2010). High-speed rail and air transport competition: Game engineering as tool for cost-benefit analysis. *Transportation Research Part B: Methodological*, *44*(7), 812–833. https://doi.org/10.1016/j.trb.2010.01.001

Albert, R., Jeong, H., & Barabási, A.-L. (2000). Error and attack tolerance of complex networks. *Nature*, *406*, 378–382. https://doi.org/https://doi.org/10.1038/35019019

Aultman-Hall, L., & Ullman, H. (2020). Long-distance and intercity travel: Who participates in global mobility? In K. G. Goulias & A. W. Davis (Eds.), *Mapping the travel behavior genome* (pp. 187–207). Elsevier. https://doi.org/https://doi.org/10.1016/B978-0-12-817340-4.00011-5

Bairras, P., & Aguas Ardaiz, I. (2025). The slow and difficult implementation of high speed rail interoperability in Europe: The case of the Atlantic Corridor. *Transportation Research Procedia*, *82*, 2454–2466. https://doi.org/10.1016/j.trpro.2024.12.204

Barabási, A.-L. (2016). *Network science*. Cambridge University Press.

Bellocchi, L., Latora, V., & Geroliminis, N. (2021). Dynamical efficiency for multimodal time-varying transportation networks. *Scientific Reports*, *11*, 23065. https://doi.org/10.1038/s41598-021-02418-5

Berche, B., von Ferber, C., Holovatch, T., & Holovatch, Y. (2009). Resilience of public transport networks against attacks. *European Physical Journal B*, *71*, 125–137. https://doi.org/10.1140/epjb/e2009-00291-3

Beria, P., Lunkar, V., Tolentino, S., Pařil, V., & Kvasnička, M. (2023). Long-distance rail in europe: Comparing the forms of head-on competition across europe. *Research in Transportation Economics*, *102*, 101367. https://doi.org/https://doi.org/10.1016/j.retrec.2023.101367

Bešinović, N. (2020). Resilience in railway transport systems: A literature review and research agenda. *Transport Reviews*, *40*, 457–478. https://doi.org/https://doi.org/10.1080/01441647.2020.1728419

Bleijenberg, A. (2020). *Air2rail: Reducing CO2 from intra-European aviation by a modal shift from air to rail* (Strategy report). Koios strategy. https://www.transportenvironment.org/wp-content/uploads/2021/07/2020_03_Air2Rail_Koios_strategy_rev.pdf

Boccaletti, S., Bianconi, G., Criado, R., del Genio, C. I., Gómez-Gardeñes, J., Romance, M., Sendiña-Nadal, I., Wang, Z., & Zanin, M. (2014). The structure and dynamics of multilayer networks. *Physics Reports*, *544*(1), 1–122. https://doi.org/10.1016/j.physrep.2014.07.001

Brandes, U. (2001). A faster algorithm for betweenness centrality. *The Journal of Mathematical Sociology*, *25*(2), 163–177. https://doi.org/https://doi.org/10.1080/0022250X.2001.9990249

Broido, A. D., & Clauset, A. (2019). Scale-free networks are rare. *Nature Communications*, *10*, 1017. https://doi.org/10.1038/s41467-019-08746-5

Brons, M., Dijkstra, L., & Poelman, H. (2023). *How fast are rail trips between eu cities and is rail faster than air?* (Tech. rep. No. WP 03/2023). Directorate-General for Regional and Urban Policy, European Commission. Luxembourg, Publications Office of the European Union. https://doi.org/10.2776/27137

Brooker, P. (2010). Fear in a handful of dust: Aviation and the Icelandic volcano. *Significance*, *7*(3), 112–115. https://doi.org/10.1111/j.1740-9713.2010.00436.x

Bruno, F., Maghrour Zefreh, M., Fröidh, O., & Cats, O. (2025). Replacing short-haul flights with train travel: Exploring impacts, capacity requirements and policy implications. *Transport Policy*, *171*, 326–343. https://doi.org/https://doi.org/10.1016/j.tranpol.2025.05.031

Büchel, B., Spanninger, T., & Corman, F. (2020). Empirical dynamics of railway delay propagation identified during the large-scale rastatt disruption. *Scientific Reports*, *10*, 18584. https://doi.org/10.1038/s41598-020-75538-z

Burghouwt, G. (2007). *Airline network development in Europe and its implications for airport planning* (1st ed.). Routledge. https://doi.org/10.4324/9781315566443

Cacchiani, V., Huisman, D., Kidd, M., Kroon, L., Toth, P., Veelenturf, L., & Wagenaar, J. (2014). An overview of recovery models and algorithms for real-time railway rescheduling. *Transportation Research Part B: Methodological*, *63*, 15–37. https://doi.org/10.1016/j.trb.2014.01.009

Calzada-Infante, L., Adenso-Díaz, B., & García Carbajal, S. (2020). Analysis of the european international railway network and passenger transfers. *Chaos, Solitons & Fractals*, *141*, 110357. https://doi.org/https://doi.org/10.1016/j.chaos.2020.110357

Cardillo, A., Zanin, M., Gómez-Gardeñes, J., Romance, M., del Amo, A. J. G., & Boccaletti, S. (2013). Modeling the multi-layer nature of the European Air Transport Network: Resilience and passengers re-scheduling under random failures. *European Physical Journal Special Topics*, *215*(1), 23–33. https://doi.org/10.1140/epjst/e2013-01712-8

Cats, O. (2016). The robustness value of public transport development plans. *Journal of Transport Geography*, *51*, 236–246. https://doi.org/10.1016/j.jtrangeo.2016.01.011

Cats, O. (2025). The long journey towards a shift to rail in the european long-distance passenger transport market. *npj Sustainable Mobility and Transport*, *2*, 7. https://doi.org/https://doi.org/10.1038/s44333-025-00025-9

Cats, O., & Jenelius, E. (2014). Dynamic vulnerability analysis of public transport networks: Mitigation effects of real-time information. *Networks and Spatial Economics*, *14*(3–4), 435–463. https://doi.org/https://doi.org/10.1007/s11067-014-9237-7

Cats, O., & Jenelius, E. (2015). Planning for the unexpected: The value of reserve capacity for public transport network robustness. *Transportation Research Part A: Policy and Practice*, *81*, 47–61. https://doi.org/10.1016/j.tra.2015.02.013

Cats, O., & Jenelius, E. (2018). Beyond a complete failure: The impact of partial capacity degradation on public transport network vulnerability. *Transportmetrica B: Transport Dynamics*, *6*(2), 77–96. https://doi.org/https://doi.org/10.1080/21680566.2016.1267596

Cats, O., Koppenol, G.-J., & Warnier, M. (2017). Robustness assessment of link capacity reduction for complex networks: Application for public transport systems. *Relia-*

bility Engineering & System Safety, *167*, 544–553. https://doi.org/10.1016/j.ress.2017.07.009

Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, *51*(4), 661–703. https://doi.org/10.1137/070710111

Clewlow, R. R., Sussman, J. M., & Balakrishnan, H. (2014). The impact of high-speed rail and low-cost carriers on European air passenger traffic. *Transport Policy*, *33*, 136–143. https://doi.org/10.1016/j.tranpol.2014.01.015

Coldren, G. M., & Koppelman, F. S. (2005). Modeling the competition among air-travel itinerary shares: GEV model development. *Transportation Research Part A: Policy and Practice*, *39*, 345–365. https://doi.org/https://doi.org/10.1016/j.tra.2004.12.001

Copernicus Emergency Management Service. (2021a). EMSR517: Flood in Germany [Activated 15 July 2021, mapping of the Bernd 2021 flood in the Ahr valley and the Erft catchment.]. https://mapping.emergency.copernicus.eu/activations/EMSR517

Copernicus Emergency Management Service. (2021b). EMSR518: Flood in Belgium [Activated 16 July 2021, mapping of the Bernd 2021 flood in the Vesdre and Meuse catchments.]. https://mapping.emergency.copernicus.eu/activations/EMSR518

Copernicus Emergency Management Service. (2021c). EMSR520: Flood in the Netherlands [Activated 17 July 2021, mapping of the Bernd 2021 flood in the Roer and Maas catchments in southern Limburg.]. https://mapping.emergency.copernicus.eu/activations/EMSR520

Cosandey, E., & Ropelius, F. (2026). Germany: Deutschlandtakt vision shapes capacity planning. *Railway Gazette International*. https://www.railwaygazette.com/passenger/2026/01/17/germany-deutschlandtakt-vision-shapes-capacity-planning/

Court of Justice of the European Union. (2009). Judgment of 19 november 2009, sturgeon and others, joined cases c-402/07 and c-432/07, ecli:eu:c:2009:716. https://curia.europa.eu/juris/liste.jsf?num=C-402/07

Curtale, R., Larsson, J., & Nässén, J. (2023). Understanding preferences for night trains and their potential to replace flights in Europe. The case of Sweden. *Tourism Management Perspectives*, *47*, 101115. https://doi.org/https://doi.org/10.1016/j.tmp.2023.101115

de Keizer, B., Kouwenhoven, M., & Hofker, F. (2015). New insights in resistance to interchange [European Transport Conference 2014]. *Transportation Research Procedia*, *8*, 72–79. https://doi.org/https://doi.org/10.1016/j.trpro.2015.06.043

de Pindray d'Ambelle, B. (2024). The impact of railway transfers on long-distance travel in europe: A stated-choice experiment on leisure travellers.

Delling, D., Pajor, T., & Werneck, R. F. (2015). Round-based public transit routing. *Transportation Science*, *49*(3), 591–604. https://doi.org/10.1287/trsc.2014.0534

Department for Transport. (2024). *TAG unit A1.3: User and provider impacts* (tech. rep.) (Transport Analysis Guidance (TAG)). Department for Transport. London. https://www.gov.uk/government/publications/tag-unit-a1-3-user-and-provider-impacts

Department for Transport (United Kingdom). (2021). Union connectivity review: Interim report. https://www.gov.uk/government/publications/union-connectivity-review-interim-report

Deutsche Bahn AG. (2025). *Integrated interim report 2025* (tech. rep.). Deutsche Bahn AG. https://zbir.deutschebahn.com/2025/en/interim-group-management-report-unaudited/quality-and-security/punctuality/

Dibbelt, J., Pajor, T., Strasser, B., & Wagner, D. (2018). Connection scan algorithm. *ACM Journal of Experimental Algorithmics*, *23*, 1.7. https://doi.org/https://doi.org/10.1145/3274661

Dijkstra, L., Poelman, H., & Veneri, P. (2019). *The EU-OECD definition of a functional urban area* (tech. rep. No. 2019/11). Organisation for Economic Co-operation and Development. Paris. https://doi.org/https://doi.org/10.1787/d58cb34d-en

Dobruszkes, F. (2011). High-speed rail and air transport competition in Western Europe: A supply-oriented perspective. *Transport Policy*, *18*(6), 870–879. https://doi.org/10.1016/j.tranpol.2011.06.002

Dobruszkes, F., Dehon, C., & Givoni, M. (2014). Does European high-speed rail affect the current level of air services? An EU-wide analysis. *Transportation Research Part A: Policy and Practice*, *69*, 461–475. https://doi.org/10.1016/j.tra.2014.09.004

Eurocontrol. (2010). *Ash-cloud of April and May 2010: Impact on air traffic* (tech. rep.). Eurocontrol Statistics and Forecast Service (STATFOR). https://www.eurocontrol.int/publication/ash-cloud-april-and-may-2010-impact-air-traffic

European Commission. (2021). *Action plan to boost long-distance and cross-border passenger rail* (COM(2021) 810). European Commission. https://transport.ec.europa.eu/news-events/news/action-plan-boost-passenger-rail-2021-12-14_en

European Commission. (2022). Commission amends ten-t proposal to reflect impacts on infrastructure of russia's war of aggression against ukraine.

European Commission. (2024, September 12). *Ten-t maps: Indicative extension to neighbouring countries* [Covers Eastern Partnership (Armenia, Azerbaijan, Georgia) and Türkiye]. Retrieved April 12, 2026, from https://transport.ec.europa.eu/transport-themes/infrastructure-and-investment/trans-european-transport-network-ten-t/tentec-information-system-and-ten-t-map-library/ten-t-maps-indicative-extension-neighbouring-countries_en

European Commission. (2025). Report from the commission to the european parliament and the council: Ninth monitoring report on the development of the rail market under article 15(4) of directive 2012/34/eu of the european parliament and of the council.

European Commission, Directorate-General for Regional and Urban Policy. (2022). Rail versus air, interactive map.

European Court of Auditors. (2018). *A European high-speed rail network: Not a reality but an ineffective patchwork* (Special Report No. 19/2018). European Court of Auditors. Luxembourg. https://op.europa.eu/webpub/eca/special-reports/high-speed-rail-19-2018/en/

European Environment Agency. (1995). *Europe's environment: The dobris assessment. chapter 3: The continent – extent and boundaries* (State of the Environment Report No. 1/1995). European Environment Agency. Copenhagen. https://www.eea.europa.eu/publications/92-826-5409-5

European Parliament and Council of the European Union. (2003). Regulation (EC) no 1059/2003 of the european parliament and of the council of 26 may 2003 on the establishment of a common classification of territorial units for statistics (NUTS).

European Parliament and Council of the European Union. (2004). Regulation (ec) no 261/2004 of the european parliament and of the council of 11 february 2004 establishing common rules on compensation and assistance to passengers in the event of denied boarding and of cancellation or long delay of flights, and repealing regulation (eec) no 295/91.

European Parliament and Council of the European Union. (2021). Regulation (eu) 2021/782 of the european parliament and of the council of 29 april 2021 on rail passengers' rights and obligations (recast).

European Parliament and Council of the European Union. (2024a). Regulation (EU) 2024/1679 of the european parliament and of the council of 13 june 2024 on union guidelines for the development of the trans-European transport network [Replaces Regulation (EU) No 1315/2013; introduces the Western Balkans, Eastern Mediterranean Core Network Corridor]. https://eur-lex.europa.eu/eli/reg/2024/1679/oj

European Parliament and Council of the European Union. (2024b). Regulation (eu) 2024/1679 of the european parliament and of the council of 13 june 2024 on union guidelines for the development of the trans-european transport network, amending regulations (eu) 2021/1153 and (eu) no 913/2010 and repealing regulation (eu) no 1315/2013. http://data.europa.eu/eli/reg/2024/1679/oj

European Union Agency for Railways. (2025). *Annual overview for interoperability* (tech. rep.). European Union Agency for Railways (ERA). Valenciennes. https://www.era.europa.eu/sites/default/files/2025-09/annual%20overview%20for%20interoperability%20-%202025.pdf

European Union and Western Balkans Six. (2017, July). Treaty establishing the transport community [Signed in Trieste on 12 July 2017; entered into force on 1 May 2019]. https://transport.ec.europa.eu/transport-themes/international-relations/enlargement/western-balkans/transport-community_en

Eurostat. (2024). Railway passenger transport statistics – quarterly and annual data. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Railway_passenger_transport_statistics_-_quarterly_and_annual_data

Eurostat. (2025a). Modal split of air, sea and inland passenger transport (`tran_hv_ms_psmod`). https://ec.europa.eu/eurostat/databrowser/view/TRAN_HV_MS_PSMOD

Eurostat. (2025b). Overview of the air passenger transport by reporting country (`avia_paoc`). https://ec.europa.eu/eurostat/databrowser/product/page/AVIA_PAOC

Eurostat. (2025c). Railway passenger transport by type of transport – national and international (`rail_pa_typepas`). https://ec.europa.eu/eurostat/databrowser/product/page/RAIL_PA_TYPEPAS

Eurostat. (2026). Air passenger transport by reporting country and main airports (national level) [Monthly air-passenger statistics; the present study uses the 2024-01 to 2026-02 window. Accessed via the Eurostat SDMX-JSON API.]. Retrieved May 19, 2026, from https://ec.europa.eu/eurostat/databrowser/view/avia_paoa

Faturechi, R., & Miller-Hooks, E. (2015). Measuring the performance of transportation infrastructure systems in disasters: A comprehensive review. *Journal of Infrastructure Systems*, *21*(1), 04014025. https://doi.org/https://doi.org/10.1061/(ASCE)IS.1943-555X.0000212

Fraszczyk, A., Lamb, T., & Marinov, M. (2016). Are railways really that bad? An evaluation of rail systems performance in Europe with a focus on passenger rail. *Transportation Research Part A: Policy and Practice*, *94*, 573–591. https://doi.org/10.1016/j.tra.2016.10.018

Frei, A., Kuhnimhof, T. G., & Axhausen, K. W. (2010). *Long distance travel in europe today: Experiences with a new survey* (tech. rep.). ETH Zurich, Institute for Transport Planning and Systems (IVT). https://doi.org/https://doi.org/10.3929/ethz-a-005976787

German Federal Ministry of Transport. (2026). The trans-european transport network (TEN-T). https://www.bmv.de/SharedDocs/EN/Articles/G/trans-european-transport-network-ten-t.html

Givoni, M., & Banister, D. (2006). Airline and railway integration. *Transport Policy*, *13*(5), 386–397. https://doi.org/https://doi.org/10.1016/j.tranpol.2006.02.001

Grolle, J., Donners, B., Annema, J. A., Duinkerken, M., & Cats, O. (2024). Service design and frequency setting for the european high-speed rail network. *Transportation Research Part A: Policy and Practice*, *179*, 103906. https://doi.org/https://doi.org/10.1016/j.tra.2023.103906

Guimerà, R., Mossa, S., Turtschi, A., & Amaral, L. A. N. (2005). The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *Proceedings of the National Academy of Sciences*, *102*(22), 7794–7799. https://doi.org/10.1073/pnas.0407994102

Hackett, A., Cellai, D., Gómez, S., Arenas, A., & Gleeson, J. P. (2016). Bond percolation on multiplex networks. *Physical Review X*, *6*(2), 021002. https://doi.org/10.1103/PhysRevX.6.021002

HTC Hanseatic Transport Consultancy and European Rail Freight Association (ERFA) and Network of European Railways (NEE) and UIRR. (2018). Economic impact assessment of the rastatt 2017 disruption (total value-added losses of EUR 2.048 billion). https://www.railfreight.com/corridors/2018/04/22/economic-damage-of-rastatt-incident-2-2-billion-euros/

International Air Transport Association. (2024). Minimum connection time (MCT) user guide. https://www.iata.org/contentassets/638f0938b3dd451b872a1d8357755421/minimum-connecting-time-user-guide_version-1.1.pdf

International Air Transport Association. (2026). *Multilateral interline traffic agreements (MITA)*. Retrieved June 28, 2026, from https://www.iata.org/en/publications/manuals/multilateral-interline-traffic-agreements/

International Railway Journal. (2026, March). ECJ rules on German track access charges. https://www.railjournal.com/policy/ecj-rules-on-german-track-access-charges/

Ippolito, N., & Cats, O. (2024). Multi-modal and multi-layer robustness analysis of the european rail and air networks. *Scientific Reports*, *14*, 26950. https://doi.org/https://doi.org/10.1038/s41598-024-76264-6

Jiao, J., Zhang, F., & Liu, J. (2020). A spatiotemporal analysis of the robustness of high-speed rail network in china. *Transportation Research Part D: Transport and Environment*, *89*, 102584. https://doi.org/https://doi.org/10.1016/j.trd.2020.102584

Junghänel, T., Bissolli, P., Daßler, J., Fleckenstein, R., Imbery, F., et al. (2021). Hydroklimatologische einordnung der stark- und dauerniederschläge in teilen Deutschlands im Zusammenhang mit dem Tiefdruckgebiet "Bernd" vom 12. bis 19. Juli 2021. *Deutscher Wetterdienst, Stand 21 July 2021*. https://www.dwd.de/DE/leistungen/besondereereignisse/niederschlag/20210721_bericht_starkniederschlaege_tief_bernd.html

Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., & Porter, M. A. (2014). Multilayer networks. *Journal of Complex Networks*, *2*(3), 203–271. https://doi.org/10.1093/comnet/cnu016

Koks, E. E., van Ginkel, K. C. H., van Marle, M. J. E., & Lemnitzer, A. (2022). Brief communication: Critical infrastructure impacts of the 2021 mid-july western european flood event. *Natural Hazards and Earth System Sciences*, *22*, 3831–3838. https://doi.org/https://doi.org/10.5194/nhess-22-3831-2022

Kreienkamp, F., Philip, S. Y., Tradowsky, J. S., Kew, S. F., Lorenz, P., et al. (2021). Rapid attribution of heavy rainfall events leading to the severe flooding in Western Europe during July 2021. *World Weather Attribution*. https://www.worldweatherattribution.org/heavy-rainfall-which-led-to-severe-flooding-in-western-europe-made-more-likely-by-climate-change/

Kristoffersson, I., & Berglund, S. (2022). Modelling connection trips to long-distance travel. *European Transport Research Review*, *14*, 30. https://doi.org/https://doi.org/10.1186/s12544-022-00556-z

Latora, V., & Marchiori, M. (2001). Efficient behavior of small-world networks. *Physical Review Letters*, *87*(19), 198701. https://doi.org/10.1103/PhysRevLett.87.198701

Law, A. M., & Kelton, W. D. (2000). *Simulation modeling and analysis* (3rd ed.). McGraw-Hill.

Li, T., & Rong, L. (2020). A comprehensive method for the robustness assessment of high-speed rail network with operation data: A case in China. *Transportation Research Part A: Policy and Practice*, *132*, 666–681. https://doi.org/10.1016/j.tra.2019.12.019

London Volcanic Ash Advisory Centre, UK Met Office. (2010). Volcanic ash advisories for the Eyjafjallajökull eruption, april–may 2010. https://www.metoffice.gov.uk/services/transport/aviation/regulated/international-aviation/vaac/va-advisories/archive

Luo, D., Cats, O., van Lint, H., & Currie, G. (2019). Integrating network science and public transport accessibility analysis for comparative assessment. *Journal of Transport Geography*, *80*, 102505. https://doi.org/https://doi.org/10.1016/j.jtrangeo.2019.102505

Lurkin, V., Garrow, L. A., Higgins, M. J., Newman, J. P., & Schyns, M. (2018). Modeling competition among airline itineraries. *Transportation Research Part A: Policy and Practice*, *113*, 157–172. https://doi.org/https://doi.org/10.1016/j.tra.2018.04.001

Malichová, E., Cornet, Y., & Hudák, M. (2022). Travellers' use and perception of travel time in long-distance trips in europe. *Travel Behaviour and Society*, *27*, 95–106. https://doi.org/https://doi.org/10.1016/j.tbs.2021.12.003

Massobrio, R., & Cats, O. (2024). Topological assessment of recoverability in public transport networks. *Communications Physics*, *7*, 108.

Mattsson, L.-G., & Jenelius, E. (2015). Vulnerability and resilience of transport systems – a discussion of recent research. *Transportation Research Part A: Policy and Practice*, *81*, 16–34. https://doi.org/https://doi.org/10.1016/j.tra.2015.06.002

Mazzocchi, M., Hansstein, F., & Ragona, M. (2010). The geo-political economy of an aborted eruption: The economic impact of the Eyjafjallajökull ash cloud. *Intereconomics*, *45*(4), 231–239. https://doi.org/10.1007/s10272-010-0339-2

Met Éireann. (2020). *Storm Ciara: Sunday 9 and monday 10 february 2020* (tech. rep.). Met Éireann. https://clidata.met.ie/cli/stormcenter/PDF/Ciara.pdf

Mohr, S., Ehmele, F., Mott, H., Aich, V., Saur, A., et al. (2023). A multi-disciplinary analysis of the exceptional flood event of July 2021 in central Europe. Part 1: Event description and analysis. *Natural Hazards and Earth System Sciences*, *23*(2), 525–551. https://doi.org/10.5194/nhess-23-525-2023

Morlotti, C., & Redondi, R. (2021). Connectivity and network robustness of European integrators. *Transportation Research Procedia*, *52*, 469–476. https://doi.org/10.1016/j.trpro.2021.01.055

Mühr, B., Eisenstein, L., Pinto, J. G., Knippertz, P., Mohr, S., & Kunz, M. (2022). *Winter storm series: Ylenia, Zeynep, Antonia (int: Dudley, Eunice, Franklin), february 2022 (NW & central europe)* (tech. rep.). CEDIM Forensic Disaster Analysis Group, Karlsruhe Institute of Technology. https://doi.org/10.5445/IR/1000143470

Nollert, M., & Niedermaier, M. (2019, November). *CODE Rhine-Alpine: New insights and challenges for action on the rhine-alpine corridor rail network* (tech. rep.). ETH Zürich, Spatial Transformation Laboratories, on behalf of the EGTC Rhine-Alpine. https://www.egtc-rhine-alpine.eu/files/2021/04/CODE_RHINE_ALPINE_EN_klein.pdf

Ortúzar, J. d. D., & Willumsen, L. G. (2011). *Modelling transport* (Fourth). John Wiley & Sons.

O'Sullivan, P. J., & Patel, T. (2004). Fragmentation in transport operations and the case for system integrity. *Transport Policy*, *11*(3), 215–225. https://doi.org/https://doi.org/10.1016/j.tranpol.2003.11.001

Otsuka, N., Günther, F. C., Tosoni, I., & Braun, C. (2017). Developing trans-European railway corridors: Lessons from the Rhine-Alpine corridor. *Case Studies on Transport Policy*, *5*(2), 351–360. https://doi.org/10.1016/j.cstp.2017.03.007

Pellegrini, P., & Rodriguez, J. (2013). Single European Sky and single European Railway Area: A system level analysis of air and rail transportation. *Transportation Research Part A: Policy and Practice*, *57*, 64–86. https://doi.org/10.1016/j.tra.2013.09.004

Pien, K.-C., Han, K., Shang, W., Majumdar, A., & Ochieng, W. (2015). Robustness analysis of the european air traffic network. *Transportmetrica A: Transport Science*, *11*(9), 772–792. https://doi.org/https://doi.org/10.1080/23249935.2015.1087233

Platform for International Rail Passenger Transport. (2024). Better railway connections for europe's passengers: A common agenda. fourth integrated progress report.

Raad voor de Leefomgeving en Infrastructuur. (2020). *Verzet de wissel: Naar beter internationaal reizigersverkeer per trein* (tech. rep.). Raad voor de Leefomgeving en In-

frastructuur (RLI). Den Haag. https://www.rli.nl/publicaties/2020/advies/verzet-de-wissel

Railteam. (2026). *Our alliance: Railteam members* [Alliance of European high-speed rail operators, founded in Brussels on 2 July 2007]. Retrieved June 28, 2026, from https://www.railteam.eu/en/why-railteam/our-alliance/

RailTech. (2025, December). FlixTrain backtracks: All 65 new high-speed trains now Germany-only, abandoning European plans. https://www.railtech.com/all/2025/12/15/flixtrain-backtracks-all-65-new-high-speed-trains-now-germany-only-abandoning-european-plans/

Šfiligoj, T., Massobrio, R., & Cats, O. (2026). Multilayer public transport networks. https://arxiv.org/abs/2603.21130

Sigmundsson, F., Hreinsdóttir, S., Hooper, A., Árnadóttir, T., Pedersen, R., Roberts, M. J., et al. (2010). Intrusion triggering of the 2010 eyjafjallajökull explosive eruption. *Nature*, *468*(7322), 426–430. https://doi.org/10.1038/nature09558

Smith, A. S. J., & Nash, C. A. (2023). Will the latest british reforms to rail passenger service procurement work? *Research in Transportation Economics*, *100*, 101321. https://doi.org/https://doi.org/10.1016/j.retrec.2023.101321

Star Alliance. (2022). Intermodal partners: Deutsche bahn (db) [Partnership effective 1 August 2022; accessed 5 July 2026].

Tanner, S., & Provoost, J. (2023). Data underlying the master thesis: Tradable mobility credits for long-distance travel in Europe, impacts on the modal split between air, rail and car [Open dataset of 112 European cities in 34 countries and 12,544 origin–destination pairs, used as the underlying data source by Ippolito and Cats (2024).]. https://doi.org/10.4121/22202389.v1

The Local. (2025). Switzerland suspends deutsche bahn trains due to chronic delays. https://www.thelocal.de/20250430/switzerland-suspends-deutsche-bahn-trains-due-to-chronic-delays

Toet, A., van Kuijk, J., Boersma, K., & Santema, S. (2026). Key factors and delivery mechanisms for passenger-oriented multimodal air journeys: A european practitioners' perspective. *Research in Transportation Business & Management*, *66*, 101610. https://doi.org/https://doi.org/10.1016/j.rtbm.2026.101610

Tomeš, Z., & Pařil, V. (2026). Night trains – sustainable alternative or niche market? *Research in Transportation Business & Management*, *64*, 101569. https://doi.org/10.1016/j.rtbm.2025.101569

Tradowsky, J. S., Philip, S. Y., Kreienkamp, F., Kew, S. F., Lorenz, P., et al. (2023). Attribution of the heavy rainfall events leading to the severe flooding in Western Europe during July 2021. *Climatic Change*, *176*(7), 90. https://doi.org/10.1007/s10584-023-03502-7

UIC. (2018). *High speed rail: Fast track to sustainable mobility* (Brochure). International Union of Railways (UIC). Paris.

UK Civil Aviation Authority. (2026). Size of UK airports [Annual passenger counts at UK commercial airports for the 2025 reference period (rundate 2026-02-25). Used to supplement the Eurostat `avia_paoa` series, which no longer covers UK airports after Brexit.]. Retrieved May 19, 2026, from https://www.caa.co.uk/data-and-analysis/uk-aviation-market/airports/uk-airport-data/

Ulfarsson, G. F., & Unger, E. A. (2011). Effects and impacts of the 2010 Eyjafjallajökull volcanic eruption in Iceland on air traffic. *Transportation Research Record*, *2214*, 144–151. https://doi.org/10.3141/2214-18

Ullman, H., & Aultman-Hall, L. (2020). Exploring motivations and barriers for long-distance trips of adult women vermonters. *Travel Behaviour and Society*, *21*, 37–47. https://doi.org/https://doi.org/10.1016/j.tbs.2020.05.007

Van Goeverden, C. D. (2009). Explaining factors for train use in european long-distance travel. *Tourism and Hospitality Planning & Development*, *6*, 21–37. https://doi.org/https://doi.org/10.1080/14790530902847038

van Goeverden, K., van Nes, R., & van Arem, B. (2019). A classification of the long-distance travel market. In F. Witlox (Ed.), *Moving towards more sustainable mobility and transport through smart systems: Proceedings of the bivec-gibet transport research days 2019* (pp. 241–255). BIVEC/GIBET.

von Ferber, C., Holovatch, T., Holovatch, Y., & Palchykov, V. (2009). Public transport networks: Empirical analysis and modeling. *The European Physical Journal B*, *68*, 261–275. https://doi.org/https://doi.org/10.1140/epjb/e2009-00090-x

Wandelt, S., & Sun, X. (2022). Lufthansa Express Rail in Germany: A critical evaluation of benefits and limitations of the intermodal network based on journey time and fares. *Multimodal Transportation*, *1*(4), 100048. https://doi.org/10.1016/j.multra.2022.100048

Wang, Z., Pei, Y., Zhang, J., & Gao, Z. (2025). Dynamic vulnerability analysis of multimodal public transport network using generalized travel costs from a multi-layer perspective. *Reliability Engineering & System Safety*, *264*, 111375. https://doi.org/10.1016/j.ress.2025.111375

Wardman, M., Chintakayala, V. P. K., & de Jong, G. (2016). Values of travel time in europe: Review and meta-analysis. *Transportation Research Part A: Policy and Practice*, *94*, 93–111. https://doi.org/https://doi.org/10.1016/j.tra.2016.08.019

Wilkinson, S. M., Dunn, S., & Ma, S. (2012). The vulnerability of the European air traffic network to spatial hazards. *Natural Hazards*, *60*(3), 1027–1036. https://doi.org/10.1007/s11069-011-9885-6

Witlox, F., Zwanikken, T., Jehee, L., Donners, B., & Veeneman, W. (2022). Changing tracks: Identifying and tackling bottlenecks in European rail passenger transport. *European Transport Research Review*, *14*(1), 7. https://doi.org/10.1186/s12544-022-00530-9

World Bank. (2020). Air transport, passengers carried. https://data.worldbank.org/indicator/IS.AIR.PSGR

Xu, Z., & Chopra, S. S. (2023). Interconnectedness enhances network resilience of multimodal public transportation systems for Safe-to-Fail urban mobility. *Nature Communications*, *14*, 4291. https://doi.org/10.1038/s41467-023-39999-w

Zhang, F., Graham, D. J., & Wong, M. S. C. (2018). Quantifying the substitutability and complementarity between high-speed rail and air transport. *Transportation Research Part A: Policy and Practice*, *118*, 191–215. https://doi.org/10.1016/j.tra.2018.08.004

Zhang, M., Huang, T., Guo, Z., & He, Z. (2022). Complex-network-based traffic network analysis and dynamics: A comprehensive review. *Physica A: Statistical Mechanics*

and its Applications, *607*, 128063. https://doi.org/https://doi.org/10.1016/j.physa.2022.128063

Zhou, Y., Wang, J., & Huang, G. Q. (2019). Efficiency and robustness of weighted air transport networks. *Transportation Research Part E: Logistics and Transportation Review*, *122*, 14–26. https://doi.org/https://doi.org/10.1016/j.tre.2018.11.008