The Global Submarine Cable Network*

Jean Imbs† Laurent Pauwels‡

February 2026

PRELIMINARY AND INCOMPLETE

Abstract

We assemble a new global dataset on fiber optic submarine cable topology, redundancy, and capacity covering 2010–2025. Applying network theory, we document rapid but unequal growth in capacity, limited gains in redundancy, and little geographic change, leaving connectivity highly concentrated. While advanced economies historically formed hub-and-spoke centers, a large group of developing economies has come to play a pivotal role in global connectivity, with densely interconnected links that weaken the traditional hub-and-spoke structure and reshape global connectivity. These emerging hubs tend to combine large populations with low political stability and are tightly linked to European and African networks, while the Americas and much of Asia remain relatively insulated. This reorganization has introduced new vulnerabilities, as increasing shares of global capacity are routed through exposed segments in international waters and along continental shelves. We discuss the forces governing the location of cable landing points.

*We would like to thank Sanjeev Goyal for very helpful conversations and seminar participants at the Academia Sinica where Imbs was visiting when this paper got started. For financial support we are grateful to the NYUAD Center for Interacting Urban Networks (CITIES), funded by Tamkeen under the NYUAD Research Institute Award CG001.

†NYUAD and CEPR

‡Stern@NYUAD

1 Introduction

Submarine fiber–optic cables carry the overwhelming majority of international data traffic and form the physical backbone of the global Internet.[1] Over the past fifteen years, this infrastructure has undergone an extraordinary technological transformation, which has generated exponential growth in the network's capabilities. Governments, regulators, and security analysts have become increasingly concerned with the resilience of this infrastructure to accidental faults and deliberate sabotage, and with the geopolitical consequences of the fact that a small number of corridors and hubs could handle a disproportionate share of global connectivity.[2]

A central contribution of this paper is the construction of a new global dataset on fiber optic submarine cables covering the period 2010–2025. The dataset aggregates information from international cable registries, operator disclosures, technical reports, maps, and manually digitized archival materials to produce a consistent panel of cables, landing points, and design capacities. A key feature of the data is its explicit geographic dimension: cables are geo-referenced at the landing-point level, and cable routes are projected onto maritime boundaries maps. We distinguish structurally integrated landing points from mere branch extensions, ensuring that the network representation correctly captures whether damage to a given point would sever connectivity along a route or merely interrupt a spur. The resulting panel spans roughly 600 active cables and 3,000 landing points in 2025, and is sufficiently rich to support extensive network analysis.

Using this dataset, we document the evolution of the global submarine cable network's topology along three layers: a binary adjacency matrix, a redundancy matrix counting parallel systems, and a capacity matrix collecting design capacities. Despite extraordinary increases in capacity, the geography of the network has remained remarkably stable, implying that network expansion has occurred primarily along the intensive margin through technological upgrading and redundancies rather than through the addition of new routes. By 2025, most countries are connected, yet the distribution of redundancy and capacity remains highly unequal.

We then examine the network's fragility. Despite being highly connected in a binary sense, the network remains fragile along both redundancy and capacity dimensions: removing links below relatively modest capacity thresholds leads to rapid fragmentation, and a large share of countries continue to rely on only one or two cables. Historically, capacity and redundancy have been concentrated in a small set of advanced economies that form the apex of a traditional hub-and-spoke structure. Since 2020, however, a new group of developing economies has emerged, accounting for a growing share of global capacity and redundancy. Largely located in the Middle East–Africa–South Asia corridor, these countries have become increasingly central to the network and, unlike traditional hubs, have developed densely interconnected links that resemble a mesh rather than a hub-and-spoke architecture. Many of

[1] Estimates routinely place the share of international traffic carried by submarine cables above 95 percent; see, for instance, Weller and Woodcock [2013], Cariolle et al. [2024].

[2] See, among others, Burnett [2021], APEC Policy Support Unit [2013], Govella [2025].

these emerging hubs are situated in environments characterized by varying degrees of political and institutional uncertainty and most of them tend to be highly populous, pointing to a tradeoff between market access and stability. They are tightly intertwined with European and African economies, in the sense that a disruption there affect primarily (and importantly) these regions. By contrast, much of Asia and the Americas remain largely insulated from disruptions in these new central economies and, more broadly, from disruptions elsewhere in the network.

The emergence of newly highly connected low- and middle-income countries with relatively low political stability raises three related questions. First, are these countries particularly vulnerable, for instance because of their geography? Second, is there anything systematic about the entities controlling these newly central connections? And third, why were these connections located there in the first place? The second part of the paper addresses these questions. We begin by showing that these new central hubs are particularly vulnerable to disruption. A large share of their cable length lies in the high seas, increasing exposure to accidental damage or sabotage, while a substantial portion of their capacity is routed through extended stretches of continental shelf: As a result, these cables spend long distances in shallow waters that are far enough from the coast to be difficult to monitor or protect.

Unsurprisingly, global fiber-optic capacity is dominated by large Western firms, both in terms of geographic reach and total controlled capacity: companies such as Google, Meta, British Telecom, or AT&T rank among the largest owners. More surprising, however, is that, outside the United States, the United Kingdom and to a lesser extent a few smaller advanced economies such as France or Singapore, the capacity associated with newly emerging central hubs is rarely controlled domestically. In countries such as Egypt or Djibouti, for instance, domestic entities account for less than five percent of total design capacity. This pattern suggests that decisions over where to locate cable landing points, and whether they function as simple branching units or full intermediate stops, are complex and typically not driven by local actors, or at least that local influence is limited. We conclude the paper with an empirical analysis of the determinants of landing-point location and operating mode. We document the importance of physical characteristics such as continental shelf configuration, near-shore sediment composition, fishing intensity, and volcanic activity, alongside economic and political factors, including a tradeoff between market access and political stability, ownership diversity, the presence of a local owner, legacy landing points, and geopolitical alignment between host countries and the countries of origin of majority owners.

Related literature Our work speaks first to the growing literature that uses submarine cable deployment as a source of variation in high–speed Internet access. A number of papers exploit the staggered arrival of new systems as a natural experiment for connectivity, and document sizable effects on employment, firm outcomes, and financial intermediation.[3] Hjort and Poulsen [2019] show that the arrival of fast Internet in African coastal regions raises employment and shifts workers into higher–skill occupations. Steinwender [2018] proposes a similar experiment based on the arrival of international telegraph cables. D'Andrea

[3]See Hjort and Tian [2024] for a review of this literature.

and Limodio [2023] study how submarine cables facilitate the adoption of real–time gross settlement systems and fintech in Africa. Jiang [2024] studies firm-level expansion in response to the development of ICT in the United States. A related line of work estimates the impact of cables on the digital divide and on Internet access prices: Cariolle [2019] shows that submarine cable deployment increases Internet penetration and reduces digital isolation in Sub–Saharan Africa, while Cariolle et al. [2024], Cariolle [2025] and O'Connor et al. [2020] document significant effects on connectivity, prices, and downstream economic activity in developing economies. These papers treat the cable infrastructure primarily as a source of quasi–experimental variation in connectivity; by contrast, we focus on the full global topology of the network and on how its multi–layer structure shapes redundancy, fragility, and the distribution of systemically important nodes. To our knowledge, the only existing study that follows a comparable strategy is Porcellacchia et al. [2025], which examines communication infrastructure over the long run, using data going back to the nineteenth century.

Recent contributions extend the analysis to broader macroeconomic outcomes and to global investment decisions. Briglauer et al. [2025] provide a cross–country assessment of the economic returns to subsea cable deployment, and Jeon and Rysman [2024] develop a structural model of data flows and cable investment on the subsea network. Caoui and Steck [2025] develop a structural model of infrastructure investment to examine whether the global internet network provides optimal diversification in suppliers.

Second, our analysis is related to the broader literature on the topology of real–world networks. Several papers have characterized the structure of the world trade network and its implications for shock propagation, using tools close to those we employ here. Kali and Reyes [2007] view international trade as a complex network and document its small–world and core–periphery properties. Fagiolo et al. [2009] study the weighted "world trade web", showing how trade intensity is distributed over heterogeneous links. A closely related literature analyzes production networks at the firm or sector level and their role in propagating shocks. Carvalho [2014] provides an early theoretical and empirical foundation for the macroeconomic relevance of input–output linkages, while Acemoglu et al. [2012] highlight how microeconomic shocks can be amplified in highly connected production networks. Several papers apply these ideas to international settings: di Giovanni and Levchenko [2010] document the importance of input–output linkages for the international transmission of shocks; Julian di Giovanni and Mejean [2018] use French firm-to-firm trade data to show how large firms shape aggregate fluctuations; Arkolakis et al. [2023] uses Chilean production network data to quantify the granular origins of aggregate fluctuations; and Bernard and Chaney [2019] show how global production networks generate concentration, market power, and amplification in international trade. Bernard et al. [2022] use customs data to uncover the sparse, hub-dominated structure of global firm-to-firm linkages. Additional theoretical and empirical contributions include Chaney [2014], who develops a dynamic model of the formation and evolution of international buyer–seller networks, and Bernard and Chaney [2019], who study how firms form global production relationships and how these shape trade flows.

In other fields, Goyal et al. [2006] describe the properties of the network of researchers in economics, using tools that are similar to those implemented in this paper. In finance,

Soramäki et al. [2007] analyze the topology of interbank payment flows, while Acemoglu et al. [2015] and Jackson and Pernoud [2021] develop theoretical and empirical frameworks for systemic risk in financial networks. Recent work in trade and urban economics studies the economic impact of transport infrastructure networks on the geography of economic activity, see for example Donaldson [2018] or Fajgelbaum and Schaal [2020]. We build on these contributions by applying similar ideas to a severely understudied but increasingly relevant physical communication infrastructure.

Finally, we connect to legal and international–relations work on the security and geopolitics of submarine cables. Burnett [2021] reviews the international law governing cable protection, emphasizing the tension between the freedom to lay and repair cables on the high seas and coastal states' rights within their exclusive economic zones. APEC Policy Support Unit [2013] quantifies the economic costs of hypothetical disruptions to major cable routes in the Pacific, stressing the importance of route diversity and redundancy. More recent contributions in security studies, such as Govella [2025], examine undersea cables as sites of geoeconomic competition and potential coercion. Cariolle [2019] also discusses how limited redundancy can expose African countries to severe digital disruptions when a single system fails. See also Clayton et al. [2025], Cannon et al. [2026], or Mohr and Trebesch [2025]. Our contribution is to provide a systematic, network–wide quantification of vulnerability using precise and recent data on the submarine cable network.

2 Data

Our analysis draws on a newly constructed global dataset of submarine cable systems, combining multiple sources to capture both the physical topology and the temporal evolution of global fiber optic digital capacity between 2010 and 2025.

2.1 Sources and coverage

The dataset merges information from the *TeleGeography Submarine Cable Map*, the *Cloudscene* registry of landing stations, the *International Cable Protection Committee*, major cable installation companies, and a set of archival operator releases and industry reports compiled manually. Each cable record includes the commissioning and decommissioning years, landing points, route length, approximate itinerary, design capacity, and ownership. Capacities are collected from further sources, including the *International Telecommunications Union* and individual cable operators or installators. Capacity is harmonized in terabits per second (Tbps). Cables are geo-localized at landing points, typically a city.

Aggregation to the country level is implemented in three distinct ways, corresponding to three network layers. In the binary network, a country pair takes the value one whenever at least one cable connects landing points located in them. The number-of-cables network records the total number of cables linking two landing points, and then aggregates these totals across all pairs of landing points belonging to the same two countries. The capacity network follows the same procedure, summing capacities first across landing-point pairs and then across country pairs. Geographic coordinates of cable landings are projected onto a

consistent coastal shapefile to allow accurate computation of over-sea distances. Although the dataset extends back to 1990, fiber-optic technology only becomes widespread in the early 2010s. Our period of analysis runs from 2010 to 2025, capturing the years during which modern fiber infrastructure developed in earnest.The resulting panel covers roughly 600 active cables and around 3,000 landing points in 2025, up from about 280 cables and 1,000 landing points in 2010.

2.2 Network construction

We represent the global submarine cable infrastructure as a time-varying weighted network $G_t = (V_t, E_t, \omega_t)$, where nodes V_t denote landing points and edges E_t link pairs of countries connected by at least one operational cable. In the analysis that follows, some results aggregate cables at the country-pair level, while others consider individual cables. Three layers of weights ω_t are used throughout the paper:

1. **Adjacency** A_t, a binary matrix with $a_{ij,t} = 1$ if at least one active cable joins i and j.
2. **Redundancy** M_t, where $m_{ij,t}$ equals the number of parallel cables between i and j.
3. **Capacity** W_t, where $w_{ij,t}$ is the total design capacity on all active cables joining i and j.

Snapshots are compiled for 2010, 2015, 2020, and 2025. Each snapshot is an undirected, unipartite graph at the level of landing points. Country-level networks are obtained by aggregating the weights M_t and W_t of all inter-country links.

The period from 2010 onward marks the effective advent of global fiber-optic connectivity. Before the late 2000s, the submarine cable infrastructure relied primarily on copper coaxial systems and early generations of analog fiber with limited amplification and bandwidth. The number of operational fiber-optic systems before 2010 was small, geographically concentrated, and technologically heterogeneous, dominated by legacy trans-Atlantic and trans-Pacific systems from the 1990s. As a result, the global network of that era was sparse and discontinuous, leaving large regions—especially Africa, South Asia, and parts of Latin America—largely unconnected or dependent on satellite backhaul.

Because of these gaps, any attempt to conduct network analysis prior to 2010 suffers from severe data incompleteness: many landing points did not exist, intercontinental routes were missing, and capacity figures were rarely disclosed or standardized. Even when available, early system capacities are several orders of magnitude smaller than those of modern systems. A technological inflection point occurred in the early 2010s with the commercial diffusion of four key technologies: Dense Wavelength-Division Multiplexing allowed a single optical fiber to carry dozens to hundreds of separate data channels simultaneously, each on a different wavelength (color) of light; Erbium-Doped Fiber Amplifiers amplified light directly without converting it to an electrical signal as earlier generations of amplifiers had to do; Coherent Transmission Technologies replaced simple "on/off" light intensity modulation with advanced digital encoding of the light wave's phase and amplitude; and Space Division Multiplexing increased capacity by adding more fibers in parallel in each cable.

The first two technologies emerged in the 1990s, the third in the 2010s, while Space Division Multiplexing is still being rolled out today. Together, these innovations resulted in extraordinary increases in capacity, from just a few tens of terabits per second (Tbps) in 2010 to several hundred Tbps by 2020. Industry announcements now point to submarine systems exceeding 1,000 Tbps as early as 2025.[4] Overall, capacity has risen exponentially over the 15-year period we study, and especially so in the most recent five years. For analytical and empirical purposes, extending the sample much earlier than 2015 offers little additional insight and risks distorting measures of density, centrality, and fragility that depend on a well-connected global backbone.

2.3 Data cleaning and missing capacity values

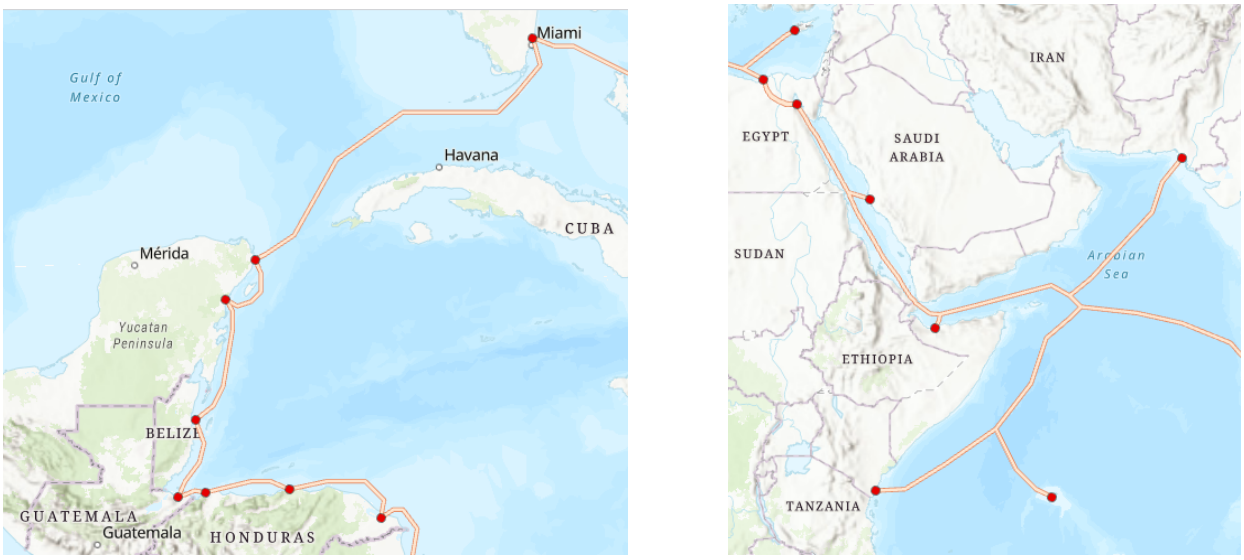
Purely domestic cables that nevertheless cross bodies of water are excluded from the network, which is designed to capture international connections only. Landing points are identified using online maps and detailed cable descriptions. Special attention is devoted to distinguishing cable extensions, i.e., segments that reach into a landing point without making that landing point an integral part of the cable, from true linkages where the landing point is a structural component of the network. In the latter case, removing the landing point or the corresponding edge severs connectivity along that cable between its upstream and downstream segments. By contrast, when a landing point merely represents an extension, its removal leaves upstream and downstream connections intact, as the core cable remains continuous beyond the destroyed landing point.

Figure 1 illustrates this distinction using snapshots of actual submarine cables in the Caribbean Sea and in the Middle East–Indian Ocean corridor. Panel (a) shows how landing points on the Yucatán coast form essential "hops": for example, they enable Miami to remain connected to Belize or Guatemala. Damage to any of these landing points would sever traffic flowing through that segment, cutting connections on both sides of the break. Panel (b), by contrast, depicts a branching configuration, in which a Kenyan landing point is not critical for maintaining a connection between Egypt and Pakistan. In such cases, traffic between those endpoints can bypass the branch entirely. Distinguishing between hop-like and branching structures is fundamental for identifying vulnerability, resilience, and overall network topology. Subject to the accuracy of the online maps we rely upon, our three network representations preserve exactly these essential differences in connectivity.

A small number of cables in the source data do not report usable information on design or lit capacity. These are typically older regional systems, legacy segments absorbed by newer consortia, or government-owned cables without public disclosure of technical specifications. Their identities and locations are listed in Figure 2. We retain their physical presence in the *adjacency* and *redundancy* layers so that connectivity is correctly measured for these layers, but assign them zero weight in the *capacity* layer. This choice prevents artificial inflation of

[4] A new cable connecting the United Kingdom, Denmark, Germany, the Netherlands, and Norway with a total capacity of 1,300 Tbps was announced in May 2024; see **??**. The press release is available at `https://www.submarinenetworks.com/en/systems/intra-europe/ioema/ioema-fibre-to-build-submarine-cable-connecting-five-northern-european-markets`.

Figure 1: The difference between hopping and branching cables



(a) Hopping Cable

(b) Branching Cable

strength for landing points connected primarily through under-documented legacy systems, while preserving the topology of the network.[5] The omission affects less than 2% of total cable count and less than 1% of observed global capacity in 2025.

All other numerical inconsistencies, duplicate landing coordinates, inactive systems listed as active, or cables lacking commissioning dates, were resolved through cross-checking with operator websites and secondary sources. The resulting dataset is internally consistent, reproducible, and suitable for both cross-sectional and time-series network analysis.

3 Topology

This section examines the evolution of the topological properties of the three network representations used in the analysis: (i) the adjacency matrix, which records whether at least one cable connects any pair of landing points; (ii) the redundancy matrix, which weighs links by the number of parallel cables between them; and (iii) the capacity matrix, which weighs links by the total available transmission capacity. For each network, we document standard measures of degree (or, in the weighted cases, strength) and distance. We report the average values, distributions, and extrema of these indicators, emphasizing how they evolve over time and what this reveals about the changing structure of the global submarine cable system.

Let N=$\{1, 2, ..., n\}$ be the set of countries in the Submarine Cable Network. For two countries $(i, j) \in N$ we characterize the existence of a cable linking the two nodes a time t with $a_{i,j,t}$ taking value 1 if a cable exists between i and j and zero otherwise. Similarly, we define the number of cables $m_{i,j,t}$ and the capacity $w_{i,j,t}$ between the two nodes. The

[5] We also assigned to those cables the average capacity of all cables commissioned in the same year, which produces virtually identical results.

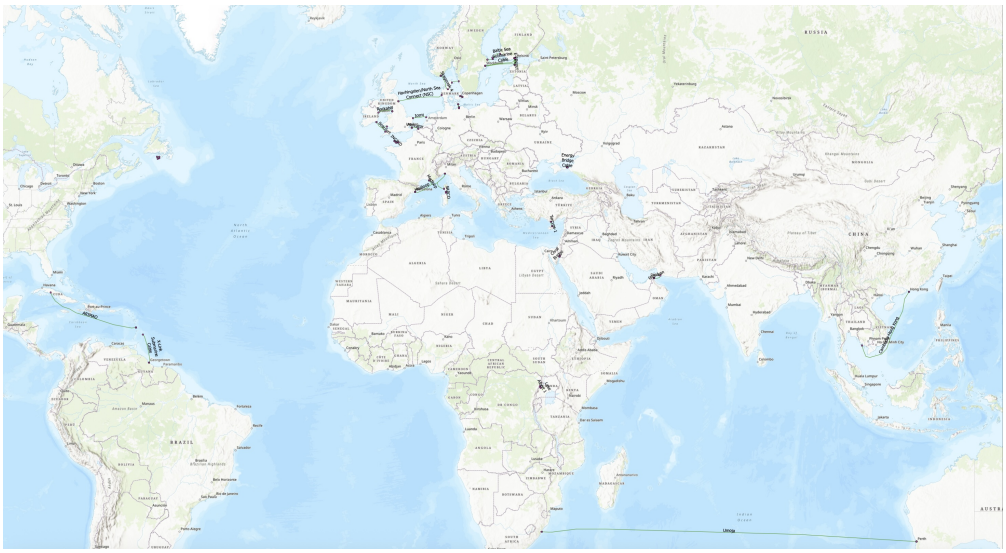


Figure 2: Cables with missing capacity data

collection of $a_{i,j,t}$ defines the adjacency network $G_t^a = (V_t, E_t, A_t)$, the collection of $m_{i,j,t}$ defines the number of cables network $G_t^m = (V_t, E_t, M_t)$ and the collection of $w_{i,j,t}$ defines the capacity network $G_t^w = (V_t, E_t, W_t)$. Let $\mathcal{N}_i(\mathbf{A}) = \{j \in N : a_{ij} = 1\}$ be the set of countries with which i has at least one cable.

Degree is defined as the number of cable connections emerging from a country i,

$$\eta_i(\mathbf{A}) = |\mathcal{N}_i(\mathbf{A})|.$$

Strength can be defined according to either weighing of the adjacency network. It can be defined as the total number of cables emerging from country i:

$$\eta_i(\mathbf{M}) = \sum_j m_{i,j}.$$

Or it can be defined as the total capacity emerging from country i:

$$\eta_i(\mathbf{W}) = \sum_j w_{i,j}.$$

All three measures can be averaged across the network, e.g.,

$$\eta(\mathbf{X}) = \frac{\sum_i \eta_i(\mathbf{X})}{n}.$$

We are interested in the cross-country dispersion in degree, its extrema, and the time pattern of those moments.

Table 1 summarizes these and several others of the main statistics discussed in Sections 3 and 4. In this Section we focus on the distributions of degree and strength across the three

Table 1: Selected topology measures

Metric	2010	2015	2020	2025
Capacity network				
Avg strength	89.15	161.22	369.94	1608.63
Std strength	158.77	219.97	660.22	2455.27
Avg weighted distance	97.35	42.98	35.90	25.73
Giant component fraction	0.80	0.90	0.97	1.00
Barrat clustering	0.63	0.67	0.68	0.71
Onnela clustering	0.04	0.06	0.01	0.02
Number of Cables network				
Avg strength	6.67	9.72	12.56	19.72
Std strength	8.36	11.29	15.01	24.14
Avg weighted distance	2.42	2.21	2.04	1.63
Giant component fraction	0.78	0.88	0.95	0.98
Barrat clustering	0.63	0.66	0.68	0.71
Onnela clustering	0.10	0.10	0.08	0.08
Adjacency (binary) network				
Avg degree	4.44	6.00	7.21	9.61
Std degree	4.86	5.85	6.56	8.77
Avg distance	3.78	3.69	3.39	3.11
Giant component fraction	0.78	0.88	0.95	0.98
Binary clustering	0.61	0.65	0.66	0.68

network layers. On average, capacity-based strength increases nearly twentyfold between 2010 and 2025, reflecting the extraordinary technological gains in fiber-optic transmission. As emphasized in Section 2.2, most of this expansion occurs in the recent period (2020–2025), coinciding with the diffusion of coherent optics and space-division multiplexing. Over the same interval, strength measured by the number of cables rises roughly threefold, signaling a broad enlargement of the network, whether through new international connections (extensive margin) or added redundancies along existing routes (intensive margin). Finally, average degree, the number of direct connections per landing point, approximately doubles between 2010 and 2025. This increase confirms some topological growth at the extensive margin, yet it remains modest compared with the dramatic surge in both cable count and total capacity.

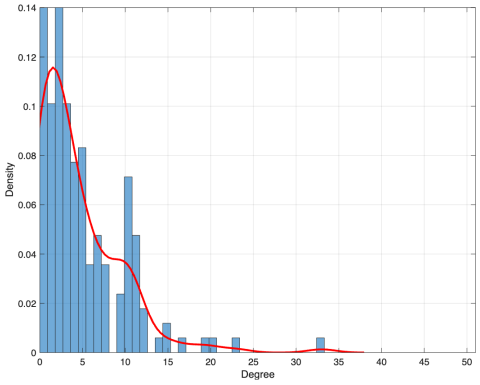
The dispersion in degree and strength rises sharply over the period. For degree and number of cables strength, the increase in dispersion closely tracks the rise in their means, leaving the coefficient of variation roughly constant, around one. This stability indicates that the topological structure of the network has changed little at the extensive margin: while most landing points have gained new connections, the relative hierarchy among nodes has remained largely intact, and the typical number of cable systems per landing point has grown proportionally.

By contrast, capacity strength dispersion grows more slowly than its mean, implying a decline in the coefficient of variation. The relative gap in capacity between large and small nodes has therefore narrowed, consistent with technological diffusion: Modern, high-capacity systems now reach a broader range of landing points rather than being concentrated in a few hubs.

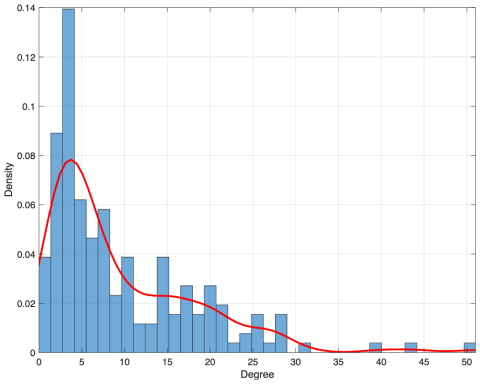
Figure 3 illustrates these developments. The density estimates for degree and for strength measured by the number of cables both shift markedly to the right between 2010 and 2025, with commensurate increases in dispersion. In contrast, the capacity-strength distributions exhibit a dramatic rightward shift accompanied by the emergence of secondary modes around 2,000–4,000 Tbps. This signals a more evenly spread of large-capacity systems and thus a smaller coefficient of variation. The figure also reveals the rise of extreme outliers, especially in the two strength measures, reflecting the emergence of nodes in the network equipped with very large capacities and redundancies.

Table 2 reports the top ten countries for both measures of strength, based on the number of cables and on total capacity, for the years 2005, 2015, and 2025. The early ranking vividly illustrates how fragmented and uneven the global fiber-optic network remained as recently as two decades ago, when smaller economies such as Estonia, Finland, or Panama appeared among the leading nodes in terms of capacity strength.

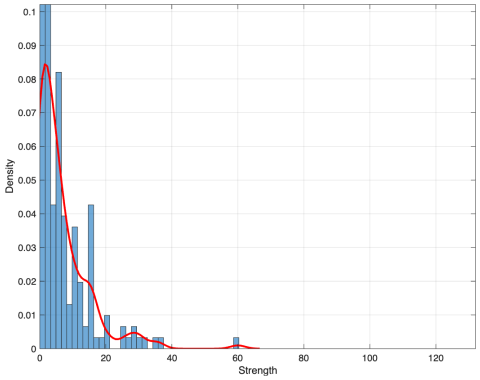
The composition of the top ten is only loosely correlated with economic size. By 2015, countries such as Egypt, South Africa, Côte d'Ivoire, and India already featured among the most connected nodes, although the rankings were still largely dominated by major economies, the United States, China, and Japan foremost among them.



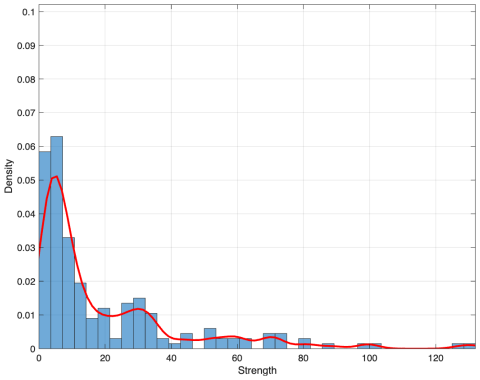
(a) Degree 2010



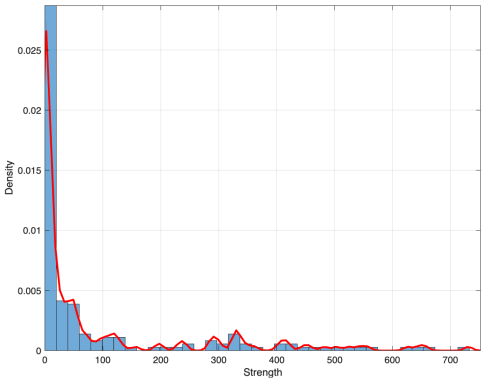
(b) Degree 2025



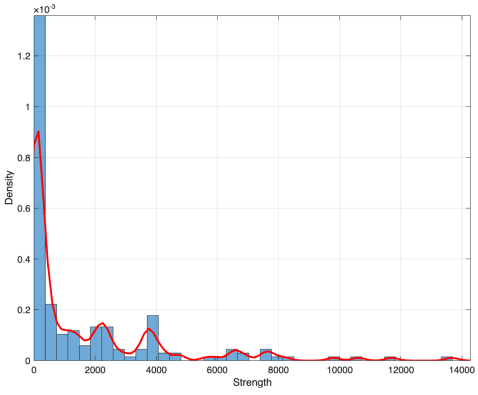
(c) Nb. cables strength 2010



(d) Nb. cables strength 2025



(e) Capacity strength 2010



(f) Capacity strength 2025

Figure 3: Degree and strength distributions across network layers (2010–2025).

By 2025, however, the landscape changes dramatically. Both rankings become dominated by emerging and developing economies: Egypt, Saudi Arabia, India, Djibouti, Oman, South Africa, the United Arab Emirates, and Pakistan all appear among the top ten in at least one strength measure. This striking shift underscores a geographical rebalancing of global connectivity and raises important questions about the resilience and geopolitical implications of a network whose critical nodes increasingly lie outside the traditional centers of economic power.

Table 2: Top-10 Countries by Strength Measures

Panel A. Number of Cables Strength

Rank	2005		2015		2025	
1	United States	42	United States	67	United States	132
2	United Kingdom	30	Egypt	47	Egypt	128
3	Japan	20	United Kingdom	46	Saudi Arabia	100
4	India	19	Saudi Arabia	42	Singapore	99
5	Singapore	19	India	41	India	89
6	Italy	17	Japan	40	United Arab Emirates	82
7	Sweden	17	South Africa	40	South Africa	80
8	China	16	China	39	Djibouti	72
9	South Korea	16	Singapore	39	Oman	72
10	Panama	15	Côte d'Ivoire	38	United Kingdom	72

Panel B. Capacity Strength

Rank	2005		2015		2025	
1	Sweden	621	Egypt	1146	United States	13095
2	United States	563	United States	1076	Egypt	11706
3	Singapore	444	Saudi Arabia	957	United Kingdom	10611
4	Denmark	441	Japan	819	Saudi Arabia	9803
5	Japan	418	United Arab Emirates	804	India	8201
6	Finland	404	India	771	Djibouti	7846
7	Estonia	404	China	738	Singapore	7674
8	South Korea	344	Singapore	680	France	7555
9	China	341	France	669	Pakistan	7522
10	Taiwan	314	Sweden	626	South Africa	6943

In the adjacency network, the distance $d(i, j, \mathbf{A})$ between countries i and j is the length of the shortest path between the two countries. $d(i, j, \mathbf{A})$ becomes infinite if there is no path between the two nodes. Since there is a maximum of $n(n-1)$ pairwise connections in the network, average distance is given by

$$d(\mathbf{A}) = \frac{\sum_i \sum_j d(i, j, \mathbf{A})}{n(n-1)}.$$

The analogue of distance in the weighted networks $\mathbf{X} = \{\mathbf{M}, \mathbf{W}\}$ introduces inverse weights $\frac{1}{m_{i,j}}$ and $\frac{1}{w_{i,j}}$, which take low values when the two nodes are linked by many or by

high capacity cables. We can define weighted distance $d(i, j, \mathbf{X})$ as the one path between i and j that minimizes $\sum_j \frac{1}{m_{i,j}}$ or $\sum_j \frac{1}{w_{i,j}}$ across all the possible paths between i and j. The average weighted distance in the network is given by

$$d(\mathbf{X}) = \frac{\sum_i \sum_j d(i, j, \mathbf{X})}{n(n-1)}.$$

Table 1 reports measures of distance across the three network layers. All distance measures decline over time, indicating a progressive tightening of the global cable system. In the binary network, average path length falls by roughly 20 percent between 2010 and 2025, pointing to incremental densification: a limited number of new edges shorten routes and create shortcuts across regions. Average weighted distance in the redundancy network falls by around 30 percent, reflecting a substantial increase in parallel systems on existing routes, consistent with a strategy of expanding redundancy and operational flexibility. The decline is far more pronounced, falling by a factor of four, when distance is weighted by capacity. This underscores the substantial improvement in link quality resulting from upgrades to high-capacity systems, made possible by exceptional technological progress over the period. It highlights the overwhelming influence of technology in shaping the evolution of the global cable network..

We draw three main conclusions from this section:

1. The physical topology of the submarine cable network has remained remarkably stable over the past two decades; what has transformed profoundly is its transmission capacity and the distribution of that capacity across nodes.

2. Capacity remains highly concentrated, with a small number of hubs accounting for a disproportionate share of the total. This inequality has narrowed somewhat in recent years, reflecting the diffusion of modern high-capacity systems.

3. By 2025, the largest capacity holders are predominantly developing and emerging economies situated along the Middle East–Africa–India corridor, marking a major geographical shift in the global distribution of digital infrastructure.

4 Fragility analysis

We next examine the sensitivity of the three cable network layers to sequential node removal. We introduce giant components, representing the largest connected subnetwork, and the percolation curve, which traces how overall connectivity deteriorates as nodes are progressively eliminated. We then turn to clustering analysis, focusing on the presence of star-like structures, i.e., nodes that combine exceptionally high strength with very low clustering coefficients, and thus play a central role in sustaining the network's global connectivity. Finally we quantify the importance of specific nodes in terms of their contribution to the overall efficiency of the network.

4.1 Giant component and percolation

A path between countries i and j exists either if $a_{i,j} = 1$, or if there is a sequence of distinct intermediary countries $j_1, j_2, ..., j_n$ such that $a_{i,j_1} = a_{j_1,j_2} = ... = a_{j_n,j} = 1$. Two countries belong to the same component if a path connects them. Components in the adjacency network $\mathbf{A}$ can be ranked by the number of countries they include. The network is said to possess a giant component, denoted $C(\mathbf{A})$, if the largest component comprises a substantial share of all countries in the network (see, for example, Goyal et al., 2006).

In weighted networks, we define the giant component as the share of total network strength accounted for by the nodes belonging to the largest component. Formally, for a weighted network $\mathbf{X}$, the giant component is given by

$$C(\mathbf{X}) = \frac{\sum_{i \in C(\mathbf{A})} \eta_i(\mathbf{X})}{\sum_i \eta_i(\mathbf{X})},$$

where $\eta_i(\mathbf{X})$ denotes the strength of node i in network $\mathbf{X}$.

Define the *percolation curve* as the share of total network strength that remains within the giant component after all links with weight below a given threshold τ are removed. Formally, let the thresholded adjacency matrix be

$$\mathbf{A}_\tau = \big[\mathbf{1}\{\mathbf{X}_{ij} > \tau\}\big],$$

and let $C(\mathbf{A}_\tau)$ denote the set of nodes belonging to the giant component of the thresholded network. The percolation curve in weighted network $\mathbf{X}$ is then defined as

$$C_\tau(\mathbf{X}) = \frac{\sum_{i \in C(\mathbf{A}_\tau)} \eta_i(\mathbf{X})}{\sum_i \eta_i(\mathbf{X})},$$

where $\eta_i(\mathbf{X})$ denotes the strength of node i in the weighted network $\mathbf{X}$. Thus, $C_\tau(\mathbf{X}) \in [0, 1]$ measures the fraction of total network strength that remains connected as links are progressively eliminated. As τ increases the network fragments and we can trace how the giant component decays. This defines a weighted percolation curve that tells us about the network's fragility (e.g., what happens if we lose all links above τTbps).

Table 1 shows that by 2020 the giant component encompasses virtually the entire submarine cable network across all three layers, capacity, number of cables, and binary connectivity. As early as 2015, approximately 90 percent of total capacity, total cables, and total linkages were already mutually connected, either directly or through intermediary nodes. This indicates an extremely dense and cohesive network, with almost no isolated landing points or sub-networks remaining. In contrast, in 2010 a number of smaller disconnected clusters still existed, representing roughly 20 percent of the network's total degree and strength. This fraction declines steadily over time, to about 10 percent in 2015, 5 percent in 2020, and only 0-2 percent by 2025—signaling the near-complete integration of the global submarine cable system.

Figure 4 displays the percolation curves for the network weighted by the number of cables (upper panel) and by total capacity (lower panel), each shown for 2010 and 2025. The upper panel illustrates the evolution of network resilience associated with the buildup of redundant links at key nodes. In 2010, the giant component collapsed once all countries with up to four cables were removed. By 2025, however, the network remains connected until the threshold reaches $\tau = 10$, indicating a substantial increase in redundancy among the most connected nodes.

Yet, the lower part of the curve reveals that redundancy remains limited for most of the network. Even in 2025, removing all countries with four cables or fewer leaves only about 20 percent of the total number of cables within the giant component. The percolation curve declines sharply at low thresholds: the connected share falls to 70 percent for $\tau = 1$ and to 43 percent for $\tau = 2$. This implies that roughly 30 percent of the world's countries are connected by a single cable, and an additional 27 percent rely on just two, underscoring the persistence of significant structural fragility in large parts of the system despite growing redundancy at its core.

The lower panel of Figure 4 displays the percolation curves for the capacity-weighted network. In 2010, nearly the entire network's capacity strength collapses once cables with capacities up to 100 Tbps are removed. Even eliminating cables below 10 Tbps cause the network to lose more than 70 percent of its total capacity—revealing a very steep percolation curve at the lower end of the capacity distribution.

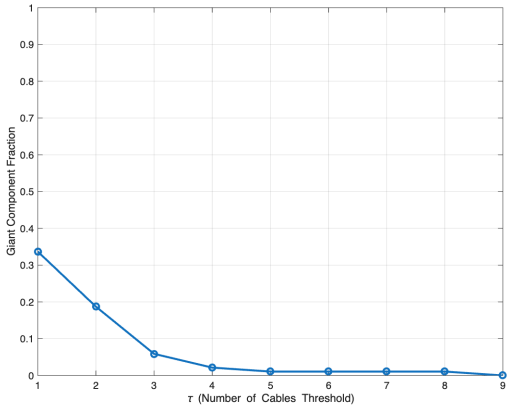
By 2025, technological progress is striking: the threshold at which the giant component collapses has increased by an order of magnitude, to around 1,000 Tbps. This reflects the diffusion of high-capacity fiber systems and the deployment of next-generation transmission technologies. Yet, the percolation curve remains extremely steep at relatively modest capacity thresholds. Removing cables with capacity below 250 Tbps still reduces the network's total capacity by almost 70 percent, and eliminating those below 500 Tbps leads to an 85 percent loss.

This precipitous decline highlights the persistent inequality in the distribution of transmission capacity across the network. Despite remarkable technological upgrading at the core, the vast majority of global connections remain supported by medium-capacity systems below 250 Tbps, leaving the overall infrastructure still highly concentrated and vulnerable to disruptions affecting the most advanced routes. This happens despite extraordinary increases in the maximum cable capacity, which by 2025 is above 4,000Tbps.

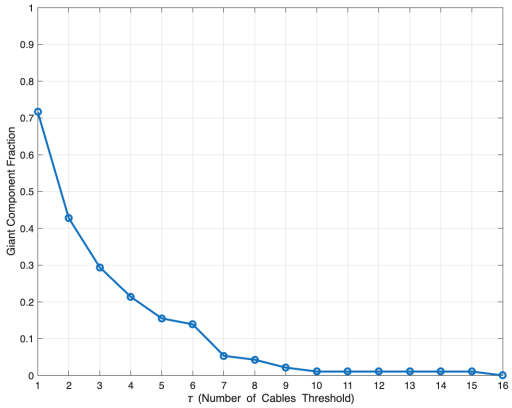
We draw two main conclusions from this section:

1. By 2025, virtually all countries are integrated into the global submarine cable system, forming a single, highly connected network with minimal geographic fragmentation.

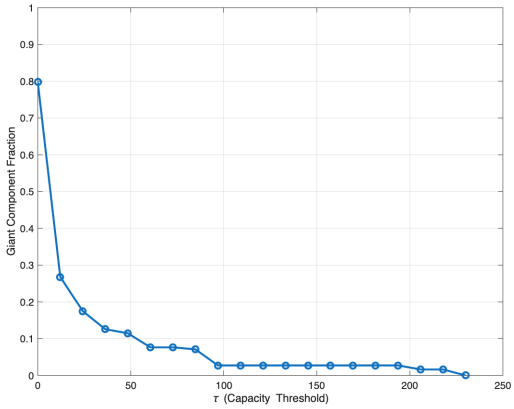
2. Despite remarkable technological advances and clear evidence of investment in redundancy, large portions of the network remain fragile, i.e., characterized by limited alternative routes and capacities still far below the technological frontier.



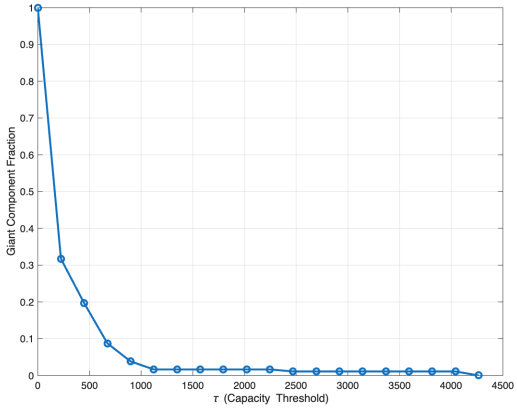
(a) Nb Cables percolation 2010



(b) Nb Cables percolation 2025



(c) Capacity percolation 2010



(d) Capacity percolation 2025

Figure 4: Percolation curves (nb cables and capacity).

4.2 Clustering

We introduce measures of clustering that capture the extent to which the countries connected with a country i are themselves connected with each other. In the adjacency network $\mathbf{A}$, the level of clustering in the neighborhood N_i of country i (that is, within the set of countries that are connected to i) is defined as

$$\kappa(\mathbf{A})_i = \frac{\sum_{l\in N_i}\sum_{k\in N_i} a_{l,k}}{\eta_i(\mathbf{A})(\eta_i(\mathbf{A})-1)},$$

for all $i \in N' \equiv \{i \in N : \eta_i(\mathbf{A}) \geq 2\}$. The denominator measures the total number of possible connections between the $\eta_i(\mathbf{A})$ connections of country i. The ratio defines the intensity of connections between country i's connections. The clustering coefficient of a network $\mathbf{A}$ is defined by the weighted average

$$\kappa(\mathbf{A}) = \frac{\sum_{i\in N'}\sum_{l\in N_i}\sum_{k\in N_i} a_{l,k}}{\sum_{j\in N'}\eta_j(\mathbf{A})(\eta_j(\mathbf{A})-1)} = \sum_{i\in N'}\frac{\eta_i(\mathbf{A})(\eta_i(\mathbf{A})-1)}{\sum_{j\in N'}\eta_j(\mathbf{A})(\eta_j(\mathbf{A})-1)}\kappa(\mathbf{A})_i.$$

The definition computes a weighted average of the clustering coefficient across all eligible countries (those with at least two connections) where each node is weighted by its share of the network's maximum possible connectivity.

There are two key measures of clustering in weighted matrices. Barrat clustering, introduced by Barrat et al (2004), measures the tendency of a node's strong connections to be mutually connected:

$$\kappa b(\mathbf{X})_i = \frac{1}{\eta_i(\mathbf{X})[\eta_i(\mathbf{A})-1]}\sum_{j\neq h}\frac{x_{ij}+x_{ih}}{2}a_{ij}a_{ih}a_{jh}$$

$\kappa b(\mathbf{X})_i$ takes high values when node i's main partners are densely connected.[6]

Onnela clustering, introduced by Onnela et al (2007), computes the geometric averages of all triangles arising from country i, accounting for weights in all three sides of the triangle:

$$\kappa o(\mathbf{X})_i = \frac{1}{\eta_i(\mathbf{A})[\eta_i(\mathbf{A})-1]}\sum_{j\neq h j,h\in N_i}(\hat{x}_{ij}\hat{x}_{jh}\hat{x}_{ih})^{\frac{1}{3}},$$

where $\hat{x}_{ij} = \{\hat{m}_{ij}, \hat{w}_{ij}\}$. The denominator performs a normalization by the total pairwise connections available among country i's connections, N_i denotes the set of countries with direct connections with i, and $\hat{x}_{ij}$ is x_{ij} normalized by its maximum value in the network. Onnela clustering evaluates whether all three edges in triangles are strong: It takes high

[6]To understand the definition of $\kappa b(\mathbf{X})_i$ consider a country i whose connections are all connected between themselves: In that case $\sum_{j\neq h}\frac{x_{ij}+x_{ih}}{2} = \eta_i(\mathbf{X})$, country i's strength. And $\sum_{j\neq h} a_{ij}a_{ih}a_{jh}$ traces the number of triangles emerging from country i that involve country j: Since there is a total of $\eta_i(\mathbf{A})$ connections emerging from i, there is at most $\eta_i(\mathbf{A})-1$ connections emerging from i and involving j. Therefore the upper bound for capacity or number of cables present in the triangles emerging from country i is given by $\eta_i(\mathbf{X})[\eta_i(\mathbf{A})-1]$, which ensures $\kappa b(\mathbf{X})_i$ takes maximum value 1. And if none of i's neighbors are connected, $\kappa b(\mathbf{X})_i = 0$.

values for well-balanced triangles, with high capacity or good redundancies. Both Onnela and Barrat clustering can be averaged across the network, simply by dividing by n.

Table 1 reports the five clustering coefficients associated with the binary, cable-weighted, and capacity-weighted network layers. In the binary network, the clustering coefficient takes high values around 0.6, so that two-thirds of the potential connections among a country's neighbors are actually present. This is consistent with the network being close to complete at the extensive margin, but of course says nothing about the redundancies or capacity available on all edges of those triangles. The Barrat clustering coefficients paint a similar picture. In both weighted networks, they take high values and increase over time, rising from about 0.63 to 0.71 between 2010 and 2025. Values in this range indicate that roughly two-thirds of a typical node's total strength participates in triangles.

By contrast, Onnela clustering coefficients are extremely low in the cable-weighted networks, consistently below 0.10. This indicates that redundancies seldom occur in triangular form: when triangles do appear, at least one side typically consists of a weak or single-cable connection. The pattern is even more pronounced in the capacity network, where Onnela clustering falls below 0.05, implying that fully interconnected triplets supported by high-capacity cables are exceedingly rare. Moreover, both coefficient decline over time, suggesting that the scarcity of strong triangles is becoming more pronounced.

These two sets of results imply that triangles are rarely balanced or complete. In most cases, two sides of the triangle (the edges linking a major hub to its neighbors) carry substantial capacity and redundancy, while the third side (the link connecting those neighbors to each other) is thin or weak. In short, large hubs sit at the apex of many high-capacity triangles, but their neighbors are only lightly connected among themselves. The fact that Barrat clustering rises over time while Onnela clustering falls slightly suggests that the hub-and-spoke configuration is intensifying over time: capacity and redundancy continue to accumulate on the hub-to-spoke connections, whereas the cross-links between spokes remain sparse and low-capacity. This implies that local resilience is deteriorating, as an increasing share of the network's connectivity and capacity becomes concentrated in already large hubs, leaving few robust alternative paths among their neighbors.

The network is thus shaped by a small number of high-capacity, highly connected hubs that enjoy substantial redundancy. We next identify these nodes more precisely. To do so, we examine the relationship between clustering measures and degree (or strength) at the country level.

Figure 5 plots, for 2025, the relationship between degree and the binary clustering coefficient. The correlation is negative and significant: countries with many connections tend to have neighbors that are only weakly interconnected. In other words, the nodes with the highest degree are precisely those whose partners do not connect to one another. They act as superstar hubs embedded in relatively sparse local neighborhoods. In the 2025 binary network, the United States is the dominant hub but also exhibits one of the lowest clustering coefficients: its connections are largely unconnected with one another. But the figure also reveals several countries with intermediate degree but high clustering, many of which lie

in the MENA–South Asia corridor. These economies have recently emerged as high-degree nodes whose links are themselves densely interconnected, thereby deviating from, and actively eroding, the traditional hub-and-spoke structure exemplified by the United States or the United Kingdom.

A similar pattern emerges in Figure 6, which shows a mildly negative correlation between Barrat clustering and strength in both weighted networks. Several of the highest-strength countries exhibit very low Barrat clustering coefficients, most notably advanced economies such as the United States, the United Kingdom, Singapore, and France (Table 2). These countries function as global hubs: they are highly connected, yet their partners are only weakly interconnected.

By contrast, clustering is substantially higher for a number of other countries that also rank highly in strength, predominantly developing economies including India, Saudi Arabia, Oman, the UAE, Pakistan, Kenya, and Djibouti. In these cases, intermediate strength, whether measured by capacity or by the number of cables, is accompanied by high or very high clustering, indicating nodes whose connections are themselves densely interconnected. Together, these patterns reveal a striking bifurcation in both weighted networks. Traditional advanced economies, led by the United States, behave as classic "superstars": highly connected nodes embedded in sparse local neighborhoods. In contrast, many economies along the Middle East–South Asia corridor are emerging as hubs characterized by both growing connectivity and increasing local cohesion, reflecting a fundamentally different mode of centrality.

These findings are reinforced by Figure 7, which plots the 2025 relationship between the Onnela clustering coefficient and strength. Developed countries with the highest strength tend to exhibit relatively low Onnela clustering, including the United States, the United Kingdom, France, and Singapore. As in Figure 6, these economies operate as global hubs within a hub-and-spoke structure. But countries with intermediate levels of capacity and redundancy display substantially higher Onnela clustering: Many are once again mostly located the Middle East–South Asia corridor. These nodes are well connected and, crucially, participate in triangles whose edges are themselves relatively strong, whether measured by cable count or capacity. This pattern characterizes India, Pakistan, Djibouti, the UAE, Saudi Arabia, and, most strikingly in the cable-weighted network, China. These economies emerge as rising "superstars," combining sizable capacity and redundancy with densely interconnected partners. Notably, virtually none of them are advanced economies.[7]

We draw two main conclusions from this section:

1. Despite dense overall connectivity, most triangles in the network contain at least one edge with limited redundancy and/or low capacity.

2. Historically, this pattern has characterized advanced economies occupying the core of a hub-and-spoke structure. More recently, however, several developing economies in the

[7]Overall levels of Onnela clustering remain extremely low: in most cases, existing triangles are highly unbalanced, with at least one edge weak in capacity or cable count. Comparisons across countries should therefore be interpreted in relative terms, as triangles remain fragile for the vast majority of nodes.

Middle East–South Asia corridor combine intermediate capacity and redundancy with densely interconnected partners, indicating a new, clustered regime of network centrality. This suggests that network resilience is increasingly concentrated in developing economies, e.g., in the Middle East–South Asia region.

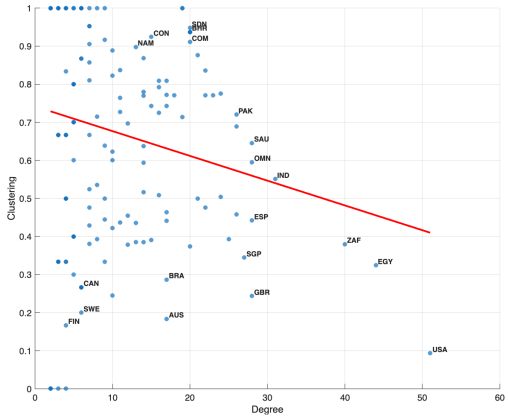
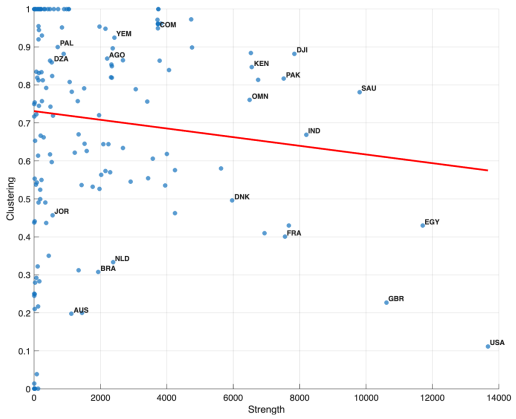
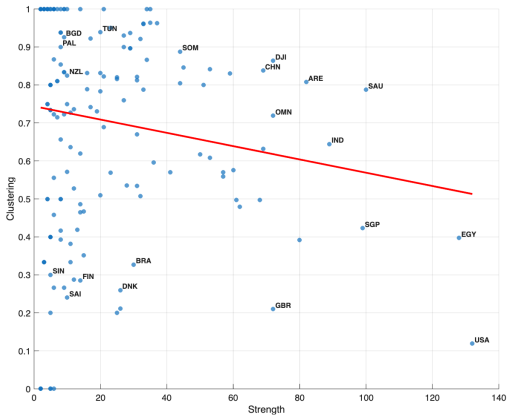


Figure 5: Binary Clustering vs. degree, r=-0.20 (p=0.01)



(a) Capacity, r=-0.10 (p=0.20)



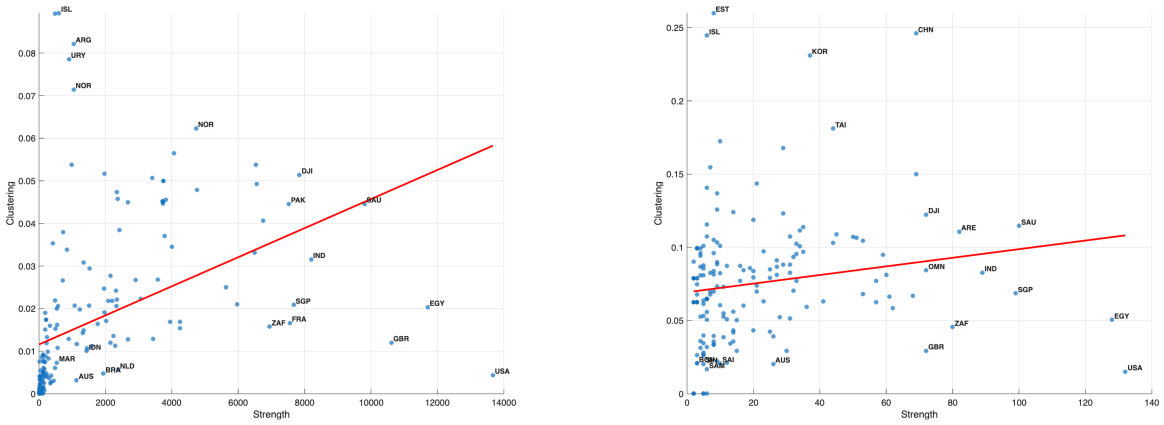
(b) Nb. of cables, r=-0.15 (p=0.05)

Figure 6: Barrat Clustering vs strength

4.3 Intercentrality

We now We propose a measure of intercentrality that captures the extent to which overall access to the network deteriorates when a given node is removed. Formally, the intercentrality of country i is defined as

$$IC_i = \sum_{j \neq k} (f^i_{jk} - f^{\bar{i}}_{jk})$$



(a) Capacity, r=0.43 (p<0.01)

(b) Nb. of cables, r=0.17 (p=0.03)

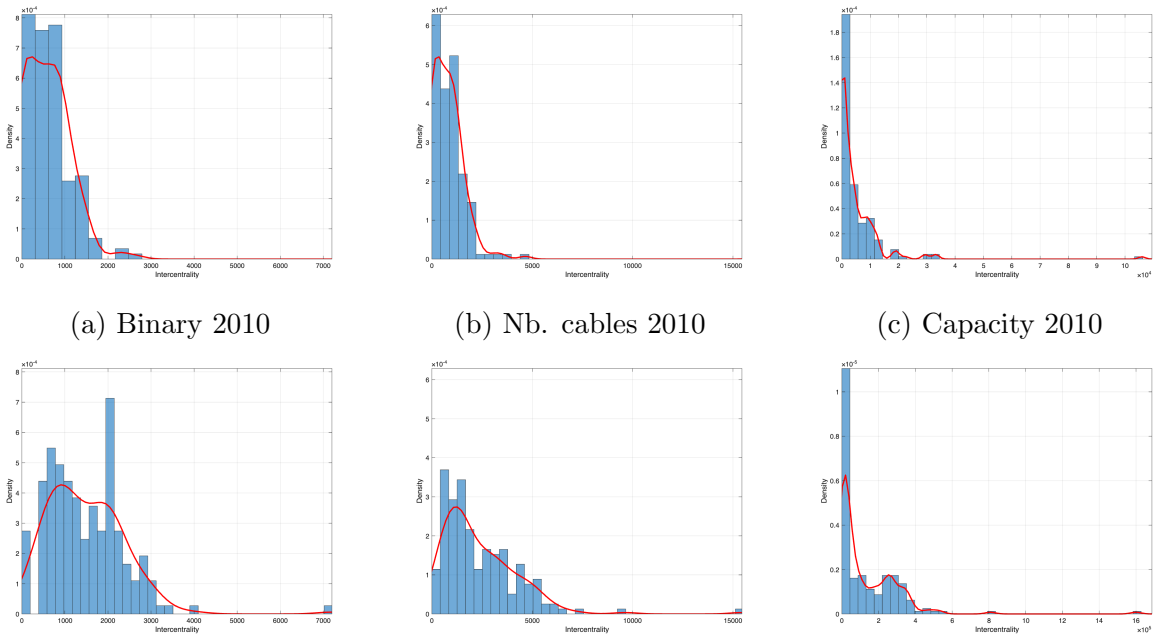
Figure 7: Onnela Clustering vs strength

where f_{jk}^{i} denotes the maximum flow value available between nodes j and k while i is present in the network and $f_{jk}^{\bar{i}}$ denotes the corresponding maximum flow when i is removed from the network. In the capacity-weighted network, for instance, IC_i measures the total loss in available capacity across all source–destination pairs resulting from the elimination of country i. We compute IC_i for each of the three network layers.

Figure 8 shows the distributions of intercentrality IC_i across the network's three layers in 2010 and 2025. All distributions are right-skewed, with a small number of countries exhibiting exceptionally high intercentrality. Notable differences, however, emerge over time. In both the binary and the number-of-cables layers, the density flattens between 2010 and 2025, and both the mean and the mode shift upward. This suggests a broadening of network centrality: more countries now account for substantial intercentrality, whether in terms of connections or in counts of physical cables. The former indicates expansion along the extensive margin, while the latter points to increasing redundancy in the network.

The distribution of capacity-based intercentrality shifts in even more dramatic fashion between 2010 and 2025.Capacity grows exponentially over the period due to rapid technological progress, making it impossible to maintain a common scale across the two distributions without sacrificing legibility. In 2010, only a small number of countries generated more than 10,000 Tbps of lost capacity were they removed from the network. By 2025, however, a far larger set of countries reaches that order of magnitude. Indeed, most countries worldwide now represent a potential loss of at least 400,000 Tbps, roughly forty times the 2010 threshold.

This expansion reflects the rise of a broad group of intermediate-strength countries, mostly developing economies from the previous section, whose removal would not be as catastrophic as that of the largest advanced economies, but would nevertheless entail losses that are one order of magnitude larger than those associated with the major economies in



(a) Binary 2010 (b) Nb. cables 2010 (c) Capacity 2010

(d) Binary 2025 (e) Nb. cables 2025 (f) Capacity 2025

Figure 8: Intercentrality distributions

2010. Despite this dispersion, the distribution remains highly skewed: a handful of countries would each generate losses exceeding 500,000 Tbps if eliminated from the network.

Table 3 reports the top 20 countries by intercentrality in terms of number of cables and capacity. In both rankings, the United States is the clear outlier, with values far exceeding those of any other country: cable intercentrality exceeds 15,000, while capacity intercentrality reaches 1.6 million Tbps. The United Kingdom is a distant second in both dimensions, with around 9,600 cables and 800,000 Tbps. From the third position onward, the rankings begin to feature countries that are neither large nor rich. Egypt ranks third in both measures. In the cable network, a substantial group of emerging and developing economies would each generate losses exceeding 4,800 cables if removed from the network. These include South Africa, Côte d'Ivoire, Saudi Arabia, India, the United Arab Emirates, Malaysia, Oman, China, Nigeria, and Brazil. Taken together, this set of countries accounts for more than 54,000 cables in potential loss.

A high number of cables does not necessarily translate into high capacity: The right panel of the table reports intercentrality measured in terms of capacity. After the United States and United Kingdom, a small set of developing countries each accounts for 300-400,000 Tbps of network capacity.[8] Taken together, the developing economies in this group account for 4.2 million Tbps, nearly twice the combined contribution of the United States and the United Kingdom. The Middle East–Africa–South Asia corridor alone represents 3.8 million

[8]Countries exceeding 400,000 Tbps include Egypt, Singapore, South Africa, and Saudi Arabia. The upper 300,000 Tbps range includes Thailand, India, Pakistan, Djibouti, the United Arab Emirates, Oman, Kenya, Somalia, and China.

Table 3: Top-20 Intercentrality Rankings in 2025

Rank	Number of Cables (2025)		Capacity (2025, Tbps)	
1	United States	15,446	United States	1,599,513
2	United Kingdom	9,641	United Kingdom	808,619
3	Egypt	7,252	Egypt	532,524
4	Italy	6,471	France	514,959
5	South Africa	6,175	Singapore	473,070
6	Singapore	6,115	South Africa	460,337
7	France	5,676	Saudi Arabia	406,318
8	Côte d'Ivoire	5,630	Thailand	376,107
9	Saudi Arabia	5,408	India	371,743
10	Portugal	5,379	Pakistan	357,464
11	Spain	5,329	Djibouti	356,129
12	India	5,213	Denmark	353,587
13	United Arab Emirates	5,119	Portugal	349,883
14	Malaysia	5,095	United Arab Emirates	349,834
15	Japan	5,087	Italy	348,255
16	Oman	5,009	Oman	341,414
17	China	5,000	Kenya	336,748
18	Nigeria	4,921	Somalia	335,794
19	Brazil	4,875	China	334,722
20	Denmark	4,817	Japan	333,918

Tbps of capacity, underscoring its growing role in anchoring global connectivity. These large capacity contributions originate in economies far from traditional centers of global activity and, in some cases, characterized by political or institutional fragility. Importantly, this group overlaps with the set of "clustered hubs," combining high strength with high clustering and forming the backbone of an emerging, qualitatively different mode of centrality.

One view is that central and intercentral nodes in the submarine cable network should primarily be advanced economies. The substantial entry of developing countries documented above challenges this perspective and raises the possibility that highly central nodes may, in fact, be politically less stable. An alternative view is that centrality reflects sheer market size rather than economic development per se. Table 4 speaks directly to these competing explanations by examining the correlates of network strength and intercentrality. The results are unambiguous. While the relationship between both measures of centrality and per capita GDP is statistically weak, central countries are not simply characterized by low political stability either. Instead, their defining feature is scale: once population is taken into account, market size dominates both income and political stability. In other words, countries are central in the submarine cable network because their large markets make them attractive connection points, even when this entails exposure to political risk. This pattern is consistent with a tradeoff between access to large consumer and production bases and the political stability of the locations through which connectivity is routed. The next question is whether this tradeoff has come at the cost of increased network fragility, by placing a disproportionate

share of capacity in politically unstable or vulnerable nodes.

Table 4: Strength and Intercentrality: Income, Political Stability, and Population

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable						
		Strength			Intercentrality	
Real per capita GDP	0.30	0.73***	0.37*	2.51*	5.51***	2.37
	(1.53)	(3.33)	(1.67)	(1.74)	(3.40)	(1.50)
Political Stability		−1.01***	0.03		−6.98***	1.99
		(−3.77)	(0.08)		(−3.56)	(0.82)
Population			0.50***			4.36***
			(4.47)			(5.42)
Constant	−1.00	−5.15**	−2.61	−8.64	−37.44**	−15.43
	(−0.51)	(−2.40)	(−1.24)	(−0.61)	(−2.37)	(−1.03)
Observations	141	141	141	141	141	141
R^2	0.02	0.11	0.22	0.02	0.10	0.26

Notes: The dependent variable is Strength in columns (1)–(3) and Intercentrality in columns (4)–(6). For legibility, Strength is divided by 1,000 and Intercentrality is divided by 10,000. Real per capita GDP and population (in millions) are measured in 2023, from the World Bank *World Development Indicators*. Both are expressed in logarithms Political Stability is measured in 2023 from the World Bank *Worldwide Governance Indicators*. All regressions are estimated by OLS. t-statistics are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We draw the following conclusions from this Section:

1. The distribution of intercentrality across countries is highly right-skewed in all three network layers and remains persistently so over time.

2. Unsurprisingly, the United States ranks as the most intercentral country. More striking, however, is the prominence further down the distribution of a large group of developing economies with intermediate levels of intercentrality, measured in both redundancy and capacity. These include the emerging "clustered hubs" located primarily in the Middle East–South Asia corridor.

3. Countries with high network centrality, both in terms of strength and intercentrality, are no longer exclusively advanced economies, but increasingly include middle-income and, in some cases, politically unstable countries. These countries tend to be highly

populous, suggesting that access to large markets has been prioritized over political stability, a tradeoff that may have introduced inherent vulnerabilities into the network.

5 The geopolitics of vulnerability

This section examines vulnerability along three dimensions. First, we document the global distribution of the losses implied by the removal of the most intercentral nodes from the network. This helps identify which regions of the world are most systematically affected when disruptions arise in key parts of the network, and, conversely, which regions remain comparatively insulated. Second, we turn to geographic sources of fragility. We assess whether the cables that are most critical for sustaining global capacity are themselves exposed to geographic risk. We do so in two ways: (i) by computing the share of each cable's length that lies in international waters, where *res communis* governance applies and vulnerability to interference is elevated; and (ii) by measuring the extent to cables remain over continental shelves, where shallow depths make them easier to access and therefore more fragile. Third, we incorporate data on cable ownership and characterize its patterns, domestic versus foreign, concentrated versus disparate, private versus state, especially for those cables identified as intercentral. This allows us to assess who controls the systems that are most consequential for the resilience of global connectivity.

5.1 Regional Interdependence

We quantify the distributional consequences across regions of removing highly intercentral nodes or cables. For each intercentral country i, we decompose the resulting loss in global capacity into the contributions of the countries that are most affected. Formally, the intercentrality of country i can be written as

$$IC_i = 2\sum_k (f^i_{1k} - f^{\bar{i}}_{1k}) + ... + 2\sum_k (f^i_{k-1,k} - f^{\bar{i}}_{k-1,k}) + 2\sum_k (f^i_{k+1,k} - f^{\bar{i}}_{k+1k}) + ... \\ + 2\sum_k (f^i_{nk} - f^{\bar{i}}_{nk}),$$

where f^i_{jk} denotes the maximum flow between j and k in the full network, and $f^{\bar{i}}_{jk}$ denotes the corresponding flow once country i is removed. Define the contribution of country j to IC_i as

$$IC_{i,j} = \frac{2\sum_k (f^i_{jk} - f^{\bar{i}}_{jk})}{IC_i}$$

The term $IC_{i,j}$ therefore measures the share of the global flow capacity lost when country i is removed that is borne by country j. In other words, it identifies which countries depend most heavily on the presence of intercentral country i for maintaining their connectivity to the global network.

Table 5 reports the allocation of capacity losses associated with the removal of each of the twenty most intercentral countries listed in Table 3. To simplify interpretation, country-level

effects are aggregated into eight regions plus Russia, with Appendix A detailing the country composition of each region. Several patterns emerge. First, for a number of countries, the largest losses occur in the region to which the removed country belongs. Removing the United States has its greatest impact on the Americas, and removing the United Kingdom or France has its largest effect on Europe. However, many cases do not follow this pattern. For Egypt, Singapore, Saudi Arabia, Thailand, India, Pakistan, Portugal, the UAE, Oman, China, and Japan, the largest losses fall outside their own region.

Second, in most of these cases the region that absorbs the largest losses is Africa. Strikingly, the removal of China or Japan generates the largest capacity loss in Africa; the same is true when removing Singapore, Thailand, India, or Pakistan. This is not a mechanical consequence of Africa's size: Africa contains 33 countries, compared to 47 in the Americas and 31 in Europe (Appendix A). Rather, it reflects the structure of the global network.

Third, after Africa, the regions most exposed to the removal of intercentral nodes are MENA and Europe. This does not arise because these regions host the bulk of the intercentral nodes: In fact, they do not, and when they do, their largest effects fall disproportionately on other regions. Instead, it reflects the underlying architecture of the network, which features dense interdependencies among Africa, MENA, and Europe. By contrast, the Americas are relatively insulated from disruptions originating elsewhere. Removing the United States naturally produces very large losses within the Americas, but removing other major hubs, including China, Japan, and large European economies, has comparatively little effect on that region.

The relative vulnerability of Africa, MENA, and Europe is further underscored by comparison with the remaining regions, whose percentage values in the table are much smaller both individually and on average. The removal of intercentral nodes has almost no effect on South Asia, Southeast Asia, East Asia, Australasia, or Russia. For Australasia and Russia, this reflects their limited integration into the global network: they have low degree and strength, and contain no intercentral nodes. More surprising is the case of Asia. The continent hosts several of the world's most intercentral hubs: Singapore, Thailand, India, Pakistan, China, and Japan all appear in the global top 20. Yet the Asian regions themselves are only weakly affected by the removal of these nodes. In other words, there is a fundamental asymmetry in the network: Africa, MENA, and Europe are highly exposed to global disruptions, while Asia and the Americas (and Australasia and Russia, though for different reasons) are largely insulated from shocks originating elsewhere in the network.

5.2 Geographic Vulnerabilities

In this section, we examine two forms of geographic vulnerability. The first concerns the legal regime governing the high seas and the risks that follow from operating submarine cables in international waters. The second relates to the extent of continental shelves surrounding landing points, and the degree of exposure they create for the cables serving each country.

The legal distinction between national and international waters is defined by the United Nations Convention on the Law of the Sea (UNCLOS). National jurisdiction extends outward

	Effect on:								
Of removing:	Africa	MENA	Europe	Americas	SE Asia	S Asia	E Asia	Australasia	Russia
US	16.9	12.1	12.4	36.9	6.9	4.9	8.3	1.7	0.0
UK	22.3	13.7	37.7	9.4	6.6	4.3	4.4	1.5	0.0
Egypt	24.2	24.1	21.4	8.3	8.7	4.5	5.0	3.8	0.1
France	21.1	23.5	26.8	9.2	7.7	4.7	4.5	2.4	0.1
Singapore	22.7	17.3	12.4	8.7	15.4	8.6	9.8	8.8	5.0
South Africa	35.2	19.9	14.4	8.6	8.2	4.4	5.8	4.3	3.5
Saudi Arabia	27.3	26.2	13.0	8.9	9.7	4.2	8.6	4.2	2.1
Thailand	23.2	18.5	13.2	9.0	15.4	8.2	8.6	9.9	2.2
India	27.2	24.5	13.4	9.2	12.0	6.5	7.2	4.4	2.1
Pakistan	29.1	25.4	13.7	9.3	10.2	6.5	5.6	4.5	2.1
Djibouti	25.9	26.5	13.7	9.3	10.2	6.4	7.8	4.5	2.2
Denmark	17.4	12.0	42.0	11.5	6.0	3.4	3.5	3.3	4.2
Portugal	38.0	17.2	13.4	11.5	7.7	4.6	4.6	4.8	2.7
UAE	28.4	24.1	13.8	9.4	10.1	5.6	7.4	4.6	2.2
Italy	22.6	22.4	24.5	10.1	8.0	5.6	4.8	4.9	2.7
Oman	24.0	23.3	14.2	9.4	9.8	7.2	7.2	4.6	5.0
Kenya	26.5	24.9	13.5	9.2	9.3	6.8	8.2	4.7	2.2
Somalia	26.1	24.8	13.5	9.2	9.3	6.8	8.2	4.7	2.2
China	23.0	17.3	13.4	9.6	18.9	5.0	10.1	2.6	0.0
Japan	22.9	17.1	13.7	11.9	13.5	5.0	10.3	3.4	2.2
Average	25.5	20.6	17.7	11.0	10.2	6.4	5.4	3.0	0.1

Table 5: Regional allocation of intercentrality losses by removed node (percent)

from the coastline across several zones. The Territorial Sea, which reaches up to 12 nautical miles from the baseline, is subject to full sovereignty of the coastal state, although foreign vessels retain a right of innocent passage, subject to certain limitations. Beyond this lies the Contiguous Zone (12–24 nautical miles), where the coastal state may enforce laws related to customs, taxation, immigration, and pollution. The outermost zone is the Exclusive Economic Zone (EEZ), extending up to 200 nautical miles; within this area, the coastal state has exclusive rights to exploit natural resources, though other states may navigate freely without engaging in economic exploitation. Waters located beyond the EEZ constitute international waters, the "high seas." These areas fall under no state's sovereignty and are governed by the principle of *res communis*: All states enjoy freedoms of navigation, overflight, scientific research, fishing, and the laying and maintenance of submarine cables. Vessels operating in the high seas are subject only to the jurisdiction of their flag state.

Submarine cables located in international waters are explicitly protected under international law. Willful or negligent damage to such cables constitutes a serious offense even when it occurs outside any national jurisdiction. Crucially, only two types of states may prosecute such offenses: the flag state of the vessel involved, or the state of nationality of the perpetrator. Unlike piracy, cable damage does not fall under universal jurisdiction: Enforcement is therefore contingent upon cooperation among states. Effective prosecution requires the ability to identify the vessel responsible, obtain adequate technical and legal evidence, and engage in diplomatic or treaty-based coordination. These constraints make the legal framework for protecting cables on the high seas straightforward in principle but difficult to enforce in practice.

We measure the share of a node's total capacity that depends on cable segments located in the high seas. The underlying motivation is that international waters are, in principle, more susceptible to deliberate cable damage. This is done by extracting the geospatial coordinates of cable routes and determining how much of the route falls inside or outside national waters.[9] We then construct bilateral measures by aggregating the relevant cable capacities linking each pair of countries, again weighting by the share of each cable that is routed through international waters. Finally, these bilateral values are aggregated to the country level to obtain a measure of the fraction of each country's total fiber-optic capacity that is located in international waters.

Figure 9 reports the distribution, across countries, of the share of total country cable capacity located in international waters. The distribution is strongly right-skewed: only a small number of countries exhibit very high values above 60 percent, while the vast majority lie below 10 percent. The mean share is 15 percent, with a standard deviation of 18.6. The upper tail is populated primarily by island economies (Cocos Islands, Bermuda, etc.) and by several African coastal countries (Angola, Senegal, Ghana, among others). In these African cases, connectivity typically relies on a small number of long offshore cables that

[9]These maps are admittedly imprecise, as the exact routes of submarine cables are not publicly disclosed. In particular, they often depict cables as straight lines between landing points. Nevertheless, the actual routes in the open ocean likely resemble these stylized paths: once cables leave the continental shelf, abyssal plains are relatively flat, burial is unnecessary, seismic risks are limited, and installation follows a cost-minimizing, and therefore distance-minimizing, logic.

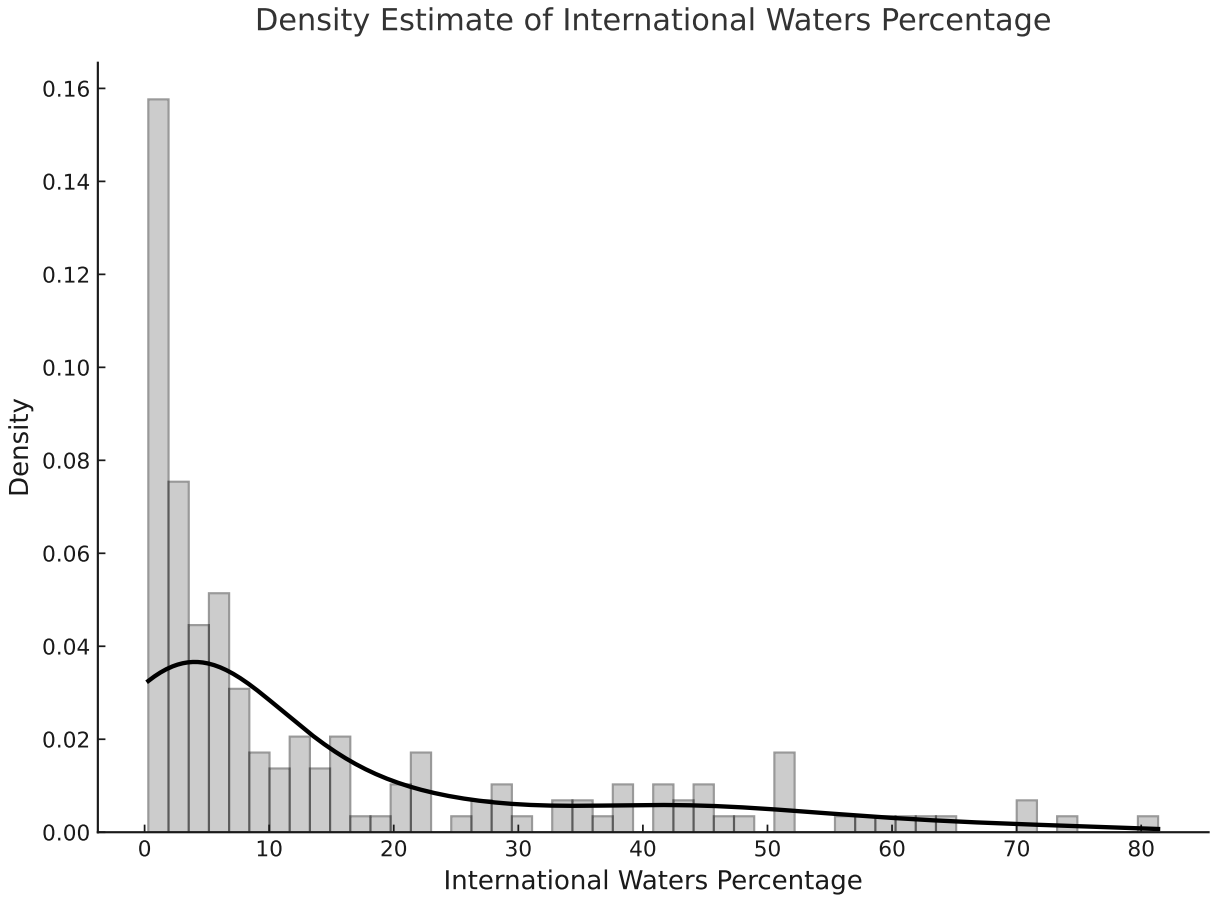


Figure 9: Density Estimate of International Waters and Distance from Shore

branch toward the coast, causing a disproportionately large share of national capacity to lie in international waters. Large countries with extensive coastlines also feature prominently at the top of the distribution. Canada ranks first, with a weighted average of 81 percent of its capacity in international waters; Spain follows with 48 percent, Brazil with 46 percent, Portugal with 45 percent, and the United States with 41 percent. Several topologically important countries also appear in the upper part of the distribution: India has 22.6 percent of its capacity in international waters, the United Kingdom 21 percent, Egypt 20 percent, Pakistan and Singapore 15 percent, and South Africa 12 percent.

Table 6: Strength, Intercentrality, and Exposure to Open Seas

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable								
	Strength				Intercentrality			
% in open seas	0.01	0.02*	0.02	0.01	0.20***	0.27***	0.25***	0.16**
	(1.46)	(1.71)	(1.63)	(0.69)	(2.85)	(3.13)	(3.11)	(2.11)
Real per capita GDP		0.31	0.74***	0.39*		2.68*	5.57***	2.71*
		(1.60)	(3.37)	(1.73)		(1.92)	(3.54)	(1.72)
Political Stability			−0.00	−0.00			−6.74***	1.31
			(−0.01)	(−0.01)			(−3.54)	(0.54)
Population				0.49***				3.95***
				(4.17)				(4.83)
Constant	1.39***	−1.42	−5.47**	−2.83	10.44***	−14.17	−41.72***	−20.27
	(6.01)	(−0.73)	(−2.55)	(−1.33)	(6.15)	(−1.03)	(−2.71)	(−1.36)
Observations	180	141	141	141	177	141	141	141
R^2	0.01	0.04	0.13	0.22	0.04	0.09	0.16	0.29

Notes: The dependent variable is Strength in columns (1)–(4) and Intercentrality in columns (5)–(8). For legibility, Strength is divided by 1,000 and Intercentrality is divided by 10,000. Real per capita GDP and population (in millions) are both measured in 2023 and expressed in logarithms, taken from the World Bank *World Development Indicators*. Political Stability is measured in 2023 from the World Bank *Worldwide Governance Indicators*. All regressions are estimated by OLS. t-statistics are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The fact that many newly emergent central countries rely on cable capacity routed through the high seas raises concerns about exposure to sabotage or accidental disruption. Table 6 examines the empirical relevance of this possibility by augmenting our analysis of network centrality with a measure of the share of each country's capacity located in international waters. The results show no systematic relationship between capacity strength and high-seas exposure: countries with large aggregate capacity are not inherently more vulnerable. By contrast, intercentral countries exhibit significantly higher shares of capacity

routed through the high seas. This pattern is not driven by high-income economies, as the (unreported) results are virtually unchanged when the sample is restricted to countries with per capita GDP below $50,000. Instead, it reflects the rise of intermediate intercentral hubs, particularly in the Middle East–Africa–South Asia corridor. The estimates indicate that the countries occupying the most influential positions in the global cable network are also those most reliant on capacity transiting international waters, where exposure to potential sabotage and disruption is greatest.

International waters are often extremely deep, typically at abyssal depths where sabotage is technically difficult. To capture a different dimension of physical exposure, we construct a second measure based on the offshore extent of the continental shelf and on the idea that waters above the shelf remain relatively shallow and therefore more vulnerable. A similar challenge to the one in international waters arises here: when damage occurs far enough offshore, identifying or prosecuting the perpetrator becomes extremely difficult. Cables located far from the coast yet still on the continental shelf face heightened exposure because their shallow depth makes them detectable with standard sensing technologies such as sonar, while their distance from shore facilitates clandestine operations.[10]

Operationally, for each landing point we identify the location at which each cable crosses the continental shelf and measure the distance from the coast to that crossing. This distance is computed separately for every cable serving a given landing point. We then construct a capacity-weighted average of these distances across all cables at the landing point, yielding a landing-point–specific measure of distance to the continental shelf. Finally, we aggregate this measure to the country level by taking a capacity-weighted average across landing points.

5.3 Ownership

We analyze the ownership structure of the key cables and nodes in the global network. We begin by documenting a set of stylized facts on major cable operators: their global reach, the capacity they control, and its concentration. We then identify the network's critical edges by computing intercentrality under cable removals and examining the cross-country distribution of capacity losses associated with those disruptions. Finally, we investigate the ownership of these critical edges to determine which firms control the infrastructure most essential to the functioning of the global network.

The left panel in Table 7 lists the operators with the widest geographic footprint, measured using the same regional classification as Table 5. The list includes the expected global players, (Google, AT&T, British Telecom, Meta, Verizon, Orange, and Vodafone) large firms headquartered in advanced economies with strong incentives to maintain worldwide connectivity. But it also features operators from emerging markets, often with close ties to their national governments (Indonesia, Malaysia, Thailand, Singapore, Vietnam, China) or large

[10]Cable installers are well aware of this vulnerability and often protect vulnerable sections by burying them. Burial, however, primarily mitigates accidental damage. It does not eliminate the risk of deliberate interference, since sonar and magnetometer systems can still detect cables in shallow or mid-depth waters, even when buried, as long as they are not at abyssal depths.

Top 20 geographic coverage		Top 20 total controlled capacity		
Company	**Regions**	**Company**	**Tbps**	**Tbps/cable**
Google — United States	8	Crosslake Fibre — Canada	4800.0	2400.0
Tata Communications — India	8	Bulk — Norway	3616.0	904.0
AT&T — United States	8	Meta — United States	3182.8	227.3
BT — United Kingdom	8	Google — United States	2974.5	135.2
Telkom Indonesia	8	euNetworks — United Kingdom	2288.0	572.0
Telekom Malaysia	8	Altibox — Norway	1920.0	1920.0
Meta — United States	8	Aqua Comms — Ireland	1726.9	287.8
Telstra — Australia	8	China Mobile — China	1457.2	112.1
Verizon — United States	8	Telecom Egypt — Egypt	805.2	67.1
Singtel — Singapore	8	Eastern Light — Sweden	768.0	384.0
National Telecom — Thailand	8	Omantel — Oman	761.0	95.1
Bharti Airtel — India	8	Orange — France	755.3	27.0
Viettel Corporation — Vietnam	8	Vodafone — United Kingdom	709.7	39.4
VNPT International — Vietnam	8	China Unicom — China	668.4	60.8
Orange — France	7	Sparkle — Italy	663.4	27.6
Vodafone — United Kingdom	7	Saudi Telecom — Saudi Arabia	660.9	73.4
China Telecom — China	7	Telxius — Spain	566.8	63.0
China Mobile — China	7	Microsoft — United States	519.6	173.2
China Unicom — China	7	Singtel — Singapore	510.2	39.2
PCCW — Hong Kong (China)	7	Reliance Jio Infocomm — India	489.0	122.3

Table 7: Top operators

private telecom groups (Tata and Bharti Airtel in India, PCCW in Hong Kong). Outside the big U.S. and Western European tech and telecom multinationals, all other globally present operators are based in Asia, likely reflecting regional geography and the pattern of Asia–rest-of-world linkages.

As shown in Section 5.1, Asia remains relatively insulated from the global network in terms of capacity. This motivates a more quantitative comparison of operators based on the capacity they control. The right panel of Table 7 provides such a comparison by listing the top 20 operators ranked by total controlled capacity. The composition of this group differs markedly from the geographic-coverage ranking, which suggests global reach does not necessarily translate into large owned capacity. Many operators from developing Asia appear among the most globally connected firms but control relatively limited capacity.

At the top of the capacity ranking are operators owning extremely high-capacity yet very short systems, such as Crosslake Fibre on Lake Ontario or Bulk and Altibox in the North and Baltic Seas, where short distance and dense fibre pairs generate unusually large terabit counts. Below these come the major tech and telecom operators from advanced economies. Meta and Google each control close to 3 Pbps, meaning that together they operate almost half of all capacity emerging from the United States (13 Pbps; see Table 2). In the United Kingdom, euNetworks (London-based but ultimately U.S.-owned) accounts for roughly 20 percent of national capacity strength, while Vodafone adds another 7 percent. In both the U.S. and the U.K., therefore, domestic ownership of submarine-cable capacity is highly concentrated, a feature that matters given their role as the two most intercentral nodes in the global network. Domestic ownership in other highly intercentral countries appears considerably less concentrated. In Egypt, Telecom Egypt controls about 7 percent of

national capacity; in France, Orange accounts for roughly 10 percent. Extreme concentration of domestic ownership thus seems to be the exception rather than the rule.

To examine this more systematically, Table 8 reports the share of capacity in the top-10 capacity-strength countries that is owned by domestic operators. The table confirms that domestic ownership is dominant only in the United States (57.6 percent) and forms a large minority in the United Kingdom (40.2 percent). In all other major capacity countries, the domestic share is much smaller, including in advanced economies such as France (12.1 percent) or Singapore (13.9 percent). The same pattern holds for the highly intercentral and capacious hubs of the MENA–Indian Ocean corridor: domestic operators control only 7 percent of Egypt's capacity, 9.5 percent of Saudi Arabia's, 13.9 percent of India's, 3 percent of Djibouti's, and 5 percent of Pakistan's. Outside the U.S. and to some extent the U.K., only a very small fraction of capacity in the most intercentral countries is owned by domestic operators. This pattern reflects the fact that most intercentral countries function primarily as transit hubs: many cables pass through their territory, boosting their available capacity and making them key hubs in the network, yet much of that capacity is not domestically owned. This is clearly the case, for example, in Djibouti, Egypt, Pakistan, or Saudi Arabia.

Country	Country Strength	Nb Operators	% Domestic Operators
United States	13095.1	17	0.576
Egypt	11706.1	1	0.069
United Kingdom	10610.7	9	0.402
Saudi Arabia	9803.3	2	0.095
India	8200.6	3	0.139
Djibouti	7846.0	1	0.031
Singapore	7674.0	7	0.139
France	7555.3	9	0.121
Pakistan	7522.2	2	0.053
South Africa	6942.9	18	0.097

Table 8: Domestic operators in top 10 capacity strength countries

6 Conclusion

This paper documents a profound reorganization of the global submarine cable network over the past fifteen years. Extraordinary growth in aggregate capacity has left the network's geography relatively unchanged, but has resulted in highly concentrated redundancy and capacity. While a small group of advanced economies historically occupied the apex of a hub-and-spoke structure, a new set of developing economies, primarily located in the Middle East–Africa–South Asia corridor, has emerged as increasingly central to global connectivity. These countries differ fundamentally from traditional hubs: their connections are densely interconnected, forming a mesh-like architecture that reshapes how shocks propagate through the network.

This reorganization has important implications for resilience. Although the network is highly connected in a binary sense, it remains fragile along redundancy and capacity dimensions. Moreover, the rise of new central hubs has shifted the locus of vulnerability. These countries tend to combine large populations with lower political stability, reflecting a tradeoff in which market access outweighs stability considerations in the placement of high-capacity infrastructure. As a result, a growing share of global capacity is routed through locations that are geographically exposed, with long stretches in international waters and extended transit along continental shelves that are difficult to monitor or protect. Control over this infrastructure is rarely domestic in these emerging hubs, underscoring that decisions over landing-point location and operating mode are typically driven by global operators rather than local actors. The empirical evidence highlights the joint importance of physical geography, economic scale, political conditions, ownership structure, and geopolitical alignment in shaping where and how cables are deployed.

Taken together, these findings suggest that the global Internet's physical backbone is increasingly anchored in locations that are both central and exposed. While this configuration reflects powerful economic incentives to serve large and growing markets, it also introduces new systemic vulnerabilities with uneven regional consequences, particularly for Europe and Africa. Understanding how these tradeoffs are managed, through redundancy, diversification, governance, and international coordination, will be critical for ensuring the resilience of global digital connectivity in the years ahead.

References

Daron Acemoglu, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012. doi: 10.3982/ECTA9623.

Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015. doi: 10.1257/aer.20130456.

APEC Policy Support Unit. Economic impact of submarine cable disruptions. Policy report, Asia-Pacific Economic Cooperation, Singapore, 2013. URL `https://www.apec.org/publications/2013/02/economic-impact-of-submarine-cable-disruptions`.

Costas Arkolakis, Federico Huneeus, and Yuhei Miyauchi. Production network formation, trade, and welfare. Technical Report 30954, National Bureau of Economic Research, 2023. URL `https://www.nber.org/papers/w30954`.

Andrew B. Bernard and Thomas Chaney. Buyer–seller networks in international trade. *Journal of International Economics*, 120:72–95, 2019. doi: 10.1016/j.jinteco.2019.05.003.

Andrew B. Bernard, Andreas Moxnes, and Yukiko Saito. Production networks, geography, and firm performance. *Journal of Political Economy*, 130(3), 2022.

Wolfgang Briglauer, Raúl L. Katz, and Juan Jung. Economic impacts of subsea cable deployment. Research paper 31, EcoAustria – Institute for Economic Research, Vienna, 2025. URL https://www.econstor.eu/handle/10419/328265.

Douglas R. Burnett. Submarine cable security and international law. *International Law Studies*, 97:1659–1682, 2021.

Brendon J. Cannon, Kazuki Matsuo, and Moeri Matsuda. Mapping undersea cable risk from bathymetry to geopolitics: Evidence-based rankings and tailored resilience strategies. *Marine Policy*, 186:107012, 2026. doi: 10.1016/j.marpol.2025.107012.

El Hadi Caoui and Andrew Steck. Content providers and the deployment of internet infrastructure. *Information Economics and Policy*, 71:101134, 2025. ISSN 0167-6245. doi: https://doi.org/10.1016/j.infoecopol.2025.101134. URL https://www.sciencedirect.com/science/article/pii/S0167624525000083.

Joël Cariolle. Telecommunication submarine-cable deployment and the digital divide in sub-Saharan Africa. Technical report, CESifo Working Paper 7415 / FERDI Working Paper 241, 2019. URL https://www.econstor.eu/handle/10419/269522.

Joël Cariolle. The impact of submarine cables on internet access price, and the role of competition and regulation. Ferdi working paper 357, FERDI, 2025. URL https://ferdi.fr/en/publications/the-impact-of-submarine-cables-on-internet-access-price-and-the-role-of-competition-a

Joël Cariolle et al. Submarine cable capacity and broadband prices. *World Development*, 172:106225, 2024. doi: 10.1016/j.worlddev.2023.106225.

Vasco M. Carvalho. From micro to macro via production networks. *Journal of Economic Perspectives*, 28(4):23–48, 2014. doi: 10.1257/jep.28.4.23.

Thomas Chaney. The network structure of international trade. *American Economic Review*, 104(11):3600–3634, 2014. doi: 10.1257/aer.104.11.3600.

Christopher Clayton, Matteo Maggiori, and Jesse Schreger. Putting economics back into geoeconomics. Mimeo, April 2025.

Angelo D'Andrea and Nicola Limodio. High-speed internet, financial technology, and banking. *Management Science*, 69(12):7277–7299, 2023. doi: 10.1287/mnsc.2023.4703.

Julian di Giovanni and Andrei A. Levchenko. The transmission of international shocks: Evidence from firm-level data. *American Economic Journal: Macroeconomics*, 2(3):95–128, 2010. doi: 10.1257/mac.2.3.95.

Dave Donaldson. Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4–5):899–934, April 2018.

Giorgio Fagiolo, Javier Reyes, and Stefano Schiavo. The evolution of the world trade web: A weighted-network analysis. *Journal of Evolutionary Economics*, 19(1):1–27, 2009. doi: 10.1007/s00191-008-0120-6.

Pablo D. Fajgelbaum and Edouard Schaal. Optimal transport networks in spatial equilibrium. *Econometrica*, 88(4):1411–1452, 2020.

Kristi Govella. Undersea cables, geoeconomics, and security in the indo-pacific. *Marine Policy*, 158:105865, 2025. doi: 10.1016/j.marpol.2024.105865.

Sanjeev Goyal, Marco J. van der Leij, and José Luis Moraga-González. Economists' networks and the diffusion of ideas. *Journal of Political Economy*, 114(2):403–412, 2006. doi: 10.1086/501171.

Jonas Hjort and Jonas Poulsen. The arrival of fast internet and employment in Africa. *American Economic Review*, 109(3):1032–1079, 2019. doi: 10.1257/aer.20161385.

Jonas Hjort and Lin Tian. The economic impact of internet connectivity in developing countries. Mimeo, August 2024.

Matthew O. Jackson and Adrien Pernoud. Systemic risk in financial networks: A survey. *Annual Review of Economics*, 13:171–202, 2021. doi: 10.1146/annurev-economics-082120-033630.

Jihye Jeon and Marc Rysman. Investment and usage of the subsea internet cable network. Technical report, Working Paper, 2024. URL `https://www.jihyejeon.com/research`. TPRC 2024 / NBER presentation.

Xian Jiang. Information and communication technology and firm geographic expansion. Working paper, December 2024.

Andrei A. Levchenko Julian di Giovanni and Isabelle Mejean. The micro origins of international business cycle comovements. *American Economic Review*, 108(1):82–108, 2018.

Raja Kali and Javier Reyes. The architecture of globalization: A network approach to international economic integration. *Journal of International Business Studies*, 38(4):595–620, 2007. doi: 10.1057/palgrave.jibs.8400286.

Cathrin Mohr and Christoph Trebesch. Geoeconomics. Kiel Working Paper 2279, Kiel Institute for the World Economy, 2025.

Alan C. O'Connor, Brian Anderson, Andrew Brower, and Sarah E. Lawrence. Economic impacts of submarine fiber optic cables and broadband connectivity in Indonesia. Working paper 0214363.202.8, RTI International, Research Triangle Park, NC, 2020. URL `https://www.rti.org/publication/economic-impacts-submarine-fiber-optic-cables-broadband-connectivity-indonesia`.

Michael Porcellacchia, Christoph Trebesch, and Benjamin Wache. Digital chokepoints: The geoeconomics of telecom networks, 1850–2025. Working paper, June 2025.

Kimmo Soramäki, Morten L. Bech, Jeffrey Arnold, Robert J. Glass, and Walter E. Beyeler. The topology of interbank payment flows. *Physica A*, 379(1):317–333, 2007. doi: 10.1016/j.physa.2006.11.093.

Claudia Steinwender. Real effects of information frictions: When the states and the kingdom became united. *American Economic Review*, 108(3):657–96, March 2018. doi: 10.1257/aer.20150681. URL https://www.aeaweb.org/articles?id=10.1257/aer.20150681.

Dennis Weller and Bill Woodcock. Internet traffic exchange: Market developments and policy challenges. *OECD Digital Economy Papers*, (207), 2013. doi: 10.1787/5k918gpt130q-en.

A Region Definitions

This section lists the individual countries included in each region introduced in Table 5

Region (Number of countries)	Countries
Americas (47)	Anguilla; Antigua and Barbuda; Argentina; Aruba; Bahamas; Barbados; Belize; Bermuda; Bonaire, Sint Eustatius and Saba; Brazil; Canada; Cayman Islands; Chile; Colombia; Costa Rica; Cuba; Dominica; Dominican Republic; Ecuador; French Guiana; Grenada; Guatemala; Guyana; Haiti; Honduras; Jamaica; Martinique; Mexico; Montserrat; Nicaragua; Panama; Peru; Saint Barthélemy; Saint Kitts and Nevis; Saint Lucia; Saint Martin; Saint Pierre and Miquelon; Saint Vincent and the Grenadines; Suriname; Trinidad and Tobago; Turks and Caicos Islands; United States; US Virgin Islands; Venezuela; Virgin Islands (U.K.); Virgin Islands (U.S.)
Africa (33)	Angola; Benin; Cameroon; Cape Verde; Comoros; Congo, Dem. Rep.; Congo, Rep.; Côte d'Ivoire; Djibouti; Equatorial Guinea; Gabon; Gambia; Ghana; Guinea; Guinea-Bissau; Kenya; Liberia; Madagascar; Mauritius; Mayotte; Mozambique; Namibia; Nigeria; Réunion; Saint Helena, Ascension and Tristan da Cunha; Sao Tome and Principe; Senegal; Seychelles; Sierra Leone; Somalia; South Africa; Tanzania; Togo; Uganda
MENA (21)	Algeria; Bahrain; Egypt; Iran; Iraq; Israel; Jordan; Kuwait; Lebanon; Libya; Mauritania; Morocco; Oman; Qatar; Saudi Arabia; Sudan; Syria; Tunisia; Turkey; United Arab Emirates; Yemen
South Asia (5)	Bangladesh; India; Maldives; Pakistan; Sri Lanka
Southeast Asia (10)	Brunei; Cambodia; Indonesia; Malaysia; Myanmar; Philippines; Singapore; Thailand; Timor-Leste; Vietnam
East Asia (4)	China; Japan; South Korea; Taiwan
Europe (31)	Albania; Belgium; Bulgaria; Croatia; Cyprus; Denmark; Estonia; Faroe Islands; Finland; France; Georgia; Germany; Gibraltar; Greece; Iceland; Ireland; Isle of Man; Italy; Jersey; Latvia; Lithuania; Monaco; Netherlands; Norway; Poland; Portugal; Romania; Spain; Sweden; Ukraine; United Kingdom
Australasia (32)	American Samoa; Australia; British Indian Ocean Territory; Christmas Island; Cocos (Keeling) Islands; Cook Islands; Fiji; French Polynesia; Guam; Kiribati; Marshall Islands; Micronesia; New Caledonia; New Zealand; Niue; Northern Mariana Islands; Palau; Papua New Guinea; Samoa; Solomon Islands; Tokelau; Tonga; Vanuatu; Wallis and Futuna; (remaining small islands included in the dataset: Papua New Guinea region groupings)
Russia (1)	Russia

Table 9: Region definitions for the 183 countries in the sample