Digital Chokepoints

Michael Porcellacchia, Christoph Trebesch, Benjamin Wache*

July 2026

Abstract

Global data flows sustain the digital economy, from AI and cloud computing to finance, yet they are routed through a network of just a few hundred physical subsea cables. The economic consequences of disrupting this critical infrastructure have never been quantified. We make three contributions to study vulnerability and strategic dependence in the global telecommunications network: (i) GlobeTel, a new dataset covering every international submarine telecom cable laid since 1850; (ii) a traffic-assignment model with endogenous congestion and rerouting; and (iii) sufficient statistics for the welfare costs of disruption requiring only observed flow shares and latency changes, validated against traceroutes, prices, and lit capacity. We find that global data flows are far more geographically concentrated than commonly understood, passing through the same maritime chokepoints that have shaped trade flows for centuries. Severing just 15 cables in the English Channel or off the US East Coast disrupts 20 to 25% of global data flows. Island economies such as Taiwan are particularly exposed, facing large potential welfare losses despite satellite access. Using our 175-year dataset, we also show that hegemons have long exploited network dominance for coercion and surveillance, and that cable investment responds systematically to geopolitical tensions.

Keywords: Geopolitical risk, sanctions, AI, technological independence, internet, digital economy

*Porcellacchia: Kiel Institute (michael.porcellacchia@ifw-kiel.de); Trebesch: Kiel Institute and CEPR (christoph.trebesch@ifw-kiel.de); Wache: CPB Netherlands Bureau for Economic Policy Analysis and Kiel Institute (b.wache@cpb.nl). We thank our discussants Anna Seim, Olaf Sleijpen, Mathias Thoenig, as well as participants at the CEPR Paris Symposium, the Geoeconomics Conference in Oxford, the 9th Annual Macroprudential Conference, the Max Planck Workshop on Geopolitical Conflict, Trade and Sanctions, the CEPR-Tübingen Conference on Geoeconomics, the HeiTüHo Workshop, and at seminars in Kiel, at the Graduate Institute Geneva, the Bundesbank, Oxford, and the European Commission. We thank Lorenzo Barrocu, Pietro Corti, Nicolo Di Luzio, Niklas Doehler, Matteo Frison, Federico Galletti, Jahana Jabbar, Sanskriti Jha, Alice Pontoglio, Mir Raisul Islam, Jonas Kaiser, Dean Sammanthan, Avi Shapiro, Joshua Sieberg, Tom Soytas, Thomas Tressel, Isabella Vornehm, Hannes Wehnhardt, Jona Wilke, Leonardo Vasquali, Yashvi Vohra, and Kexuan Zhou for excellent research assistance. Christoph Trebesch acknowledges financial support from the European Union (ERC, Great.Power.Finance, 101087838). Views and opinions are those of the authors and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

1 Introduction

Modern economies depend on the uninterrupted flow of digital information within and across borders. The subsea cable network makes this possible. Almost all intercontinental data traffic travels through a network of about 500 submarine cables, making it a critical infrastructure of globalization. These cables sustain cloud computing, AI, finance, logistics, and international trade worldwide. Historically, cables have played a similarly critical role, as telegraphy enabled the rapid transmission of news, military orders, or financial transactions (Steinwender, 2018). A growing economics literature shows that such digital connectivity generates substantial welfare gains through lower transaction costs, higher productivity, and deeper international integration (Farboodi and Veldkamp, 2020; Brynjolfsson et al., 2025; Veldkamp and Chung, 2024; Hjort and Tian, 2025). Yet, existing research has almost exclusively studied the *gains* from connectivity and data access.

Much less is known about the reverse margin: how exposed is the global economy to large-scale disruptions of the infrastructure carrying global data flows? How vulnerable is digital globalization to sabotage, coercion, and geopolitical fragmentation? These questions are becoming increasingly salient as global communication networks are both physically vulnerable and geopolitically exposed. Submarine cables are concentrated in a small number of geographically fixed systems and have repeatedly become targets of espionage, coercion, and sabotage, as illustrated by recent incidents in the Baltic Sea, the Red Sea, and around Taiwan.

This paper is the first quantitative analysis of vulnerability and strategic dependence in global digital infrastructure, spanning 150 years. We study two forms of disruption. The first is physical, when coordinated cable cuts or military conflict disable critical parts of the network. The second operates through sanctions and embargoes, which exclude a country from cables controlled by a hostile power. We do not study isolated cable outages, which are frequent and usually manageable. Our focus is on rare, large-scale events in which substantial parts of a country's network become unavailable simultaneously.

Conceptually, subsea telecom cables are a prototypical chokepoint, as defined in the emerging literature on geoeconomics (see Farrell and Newman, 2019; Clayton et al., 2026, 2025a; Mohr and Trebesch, 2025). In this literature, a chokepoint is a node in a network where the target faces high dependence, influence and/or control by an external actor, and limited substitutability. Submarine cables fit this definition closely. First, economies worldwide crucially depend on data flows, but these flows are concentrated on a limited number of cables. Globally, there are a few hundred cables, but individual countries depend on far fewer. As of 2023, South Korea was connected to the world through just 10 fiber-optic cables, Taiwan through 14, Japan through 28, and Great Britain through 60. Second, a few large countries and firms control a disproportionate share of this infrastructure, giving them high leverage for exclusion, surveillance, and sanctions against countries they seek to target.

Third, disruption is cheap while defense is expensive: a fishing vessel dragging an anchor is enough to cut a cable, while continuously surveilling thousands of kilometers of seabed is very costly. Repair capacity is equally constrained. As of 2025, there were only about 70 cable laying ships capable of repairing subsea cables, and just about 20 of them specialize in repairs (Subramanian, 2025; International Cable Protection Committee, 2025). In active conflict zones, such as the war theater in the Strait of Hormuz, repairs can be impossible altogether.

Finally, there are few viable alternatives to cables. Satellites are often seen as a key alternative, but our data show they are a poor substitute due to their limited bandwidth and technical properties. A single modern subsea cable can transmit more than 500 terabits per second. This is more than the combined capacity of *all* of the 7,000 SpaceX satellites in orbit as of 2025. Satellites also face several technical constraints: they perform poorly in cloudy conditions, are vulnerable to electronic jamming, have higher latency (transmission times), and are more expensive to build and maintain. Their main comparative advantage lies in broad coverage, particularly in remote areas lacking cable infrastructure or reliable mobile networks.

To quantify systemic risks in global data flows across space and time, we construct the Global Telecommunication Database, or *GlobeTel*, covering every known international submarine telecommunication cable (telegraph, telephone, and internet) from 1850 to 2025, as well as every satellite launched to space. The dataset combines information on geography, ownership, capacity, and technical characteristics, allowing us to reconstruct the evolution of the global telecommunications network over more than 150 years.

We combine these data with a quantitative model in which digital connectivity enters production as an intermediate input. Firms draw knowledge from geographically dispersed sources, and the cost of reaching a source is the latency of the route that carries the data. Latency is not a primitive of the model but an equilibrium object: competitive carriers assign each flow to the cheapest available path, the resulting traffic congests that path, and demand, routing, and delay are determined jointly. A cable cut or an embargo therefore imposes costs beyond the traffic directly displaced, since re-routed data travel longer paths and the congestion they generate slows the remaining flows. This yields a welfare cost for any restriction of the network. We define a country's *vulnerability* as the welfare cost of cutting a set of cables, and its *dependence* on another country as the welfare cost of being denied that country's cables. Both are computed by perturbing the network and re-solving for equilibrium, so the two measures derive from a single framework rather than from separate indicators.

Estimating the model amounts to reconstructing the physical and economic architecture of the global internet. This requires recovering several key objects that are not directly observed, including the routes submarine cables follow on the ocean floor, the international terrestrial backbone connecting them, cable-level traffic, and the equilibrium routing of global data flows. We reconstruct these objects with data from 2023, by combining information on cable geography, ownership, capacity, bathymetry, and observed bilateral traffic, together with a traffic-assignment model featuring endogenous congestion and rerouting. The result is a quantitative representation of the global telecommunications network that can be used to trace how data move across individual cables and how disruptions propagate through the system.

The resulting framework generates a rich set of predictions that are not targeted in estimation. We validate the reconstructed network against observed cable lengths and internet latencies, compare predicted routes to traceroute data, test the implied relationship between latency and transport prices using leased-line quotes, and compare predicted cable utilization to observed lit capacity. Traceroute evidence further confirms the model's prediction that satellites are used almost exclusively for first- and last-mile connections rather than as substitutes for submarine cables. Together, these validation exercises show that the model provides a realistic representation of how global data move through today's global cable network.

Our quantitative framework is calibrated to the modern internet, which offers unusually rich data. Yet today's global network has developed during an exceptionally stable geopolitical period. Post-1990 data therefore provide limited evidence on how wars, sanctions, and geopolitical rivalry reshape global communications and telecom infrastructure. History provides the missing variation. We exploit it by compiling the first systematic evidence on "digital blockades" and the evolution of the global telecommunications network over the past 175 years.

History shows that states have repeatedly weaponized telecom networks during periods of geopolitical rivalry. They cut submarine cables during war, excluded rivals from cable networks, and exploited network control for coercion and large-scale espionage. We map, for the first time, how the United States severed Spain's cables to Cuba during the Spanish-American War of 1898 and how Britain cut Germany's overseas cables in 1914, largely isolating Germany from intercontinental communications throughout World War I. Our evidence shows that such disruptions need not be temporary. Countries can remain partially or fully disconnected from global communications for years or even decades. Finally, we show that geopolitical tensions systematically reshape the geography of global telecom investment.

We establish three main findings:

1. **Hegemons dominate global telecom networks, giving them considerable leverage for surveillance and coercion.** Dominant powers have controlled the international telecom network throughout its history: the British Empire from the 1860s, and the United States after 1945, especially after the fiber-optic revolution of the 1980s. Hegemons act as global backbone providers, so their dominance exceeds what ownership shares alone suggest. We estimate that roughly 70% of global data flows today traverse at least one US-owned cable somewhere along their route, up from less than 20% in the early 1990s. Historically, more than half of international telegraph traffic passed through British-controlled cables.

   This dominance yields both surveillance capacity and coercive leverage. A hegemon can intercept the data routed over its cables, so the share of global flows traversing them is a measure of its surveillance exposure. It can also deny rivals access, and we estimate the welfare losses of this type of digital sanction based on our model and 2023 data. A simulated US cable embargo generates substantial welfare losses, which are largest in Central and South America, where substitutes for US-owned infrastructure are scarcest. China's leverage over its Asian neighbours is smaller, since its network position is less central and its targets retain more alternatives. Surveillance exposure is global in scope, while coercive leverage is more regionally concentrated.

2. **Global data flows are surprisingly concentrated in a few vulnerable corridors.** One of the paper's central insights is that the digital economy

is vulnerable to disruption in a small number of physical corridors, or ”digital chokepoints”. The modern cable network looks dense, but we find traffic to be heavily concentrated on a handful of routes with limited redundancy. Our framework of traffic routing in equilibrium makes this concentration visible in a way that a plain cable map cannot. The picture that emerges runs counter to the image of a placeless ”world wide web” or ”cloud”. Due to geography, international data flows are largely funneled through the same maritime passages that have defined the movement of goods and naval power for centuries, including Suez, Malacca, Gibraltar, and the English Channel. We also find a striking concentration off the US East Coast, where a small number of landing stations account for a large share of transatlantic traffic. The geography of digital globalization, in other words, is less virtual than commonly assumed.

We then quantify the vulnerability this creates. Removing any single cable is close to costless, since operators can reroute around individual failures easily. Risks are much higher in case of coordinated cable cuts in critical corridors, where many cables run together. The most exposed two-degree grid cell carries 8.2% of global traffic on its own, and we estimate that cutting just 15 cables in the English Channel or off the US East Coast (NY–NJ–Virginia shelf) would disrupt 20 to 25% of global data flows. The rapid advance of unmanned underwater drones makes such an attack an increasingly plausible scenario. Taiwan is an illustrative case. We simulate a digital blockade in which all 14 of its international cables are cut and find severe welfare losses that satellite connectivity cannot offset, even under the generous assumption that Taiwan has access to Starlink, which it currently lacks. Once cables are severed, no alternative routes remain, and satellite capacity is orders of magnitude too small to substitute. More broadly, we find that island and remote economies bear the largest risk.

3. **Geopolitical tensions drive cable investment and decoupling.** Using historical and modern data, we show that cable investment responds to geopolitical shocks. We start with a case study of the weaponization of global telegraphy during 1898–99. In the years that followed, France, Germany, and the United States considerably expanded their own networks, often in the same corridors in which Britain had already built substantial capacity. This pattern also appears in corridor-level regressions, using our full 1850–2025 dataset. We find that cable building rises significantly in

periods with geopolitical tensions, and the response is strongest among countries on poor terms with the hegemon. Even the ongoing investment wave follows the same logic. We collect new data on cable projects planned up to 2035 and find that several of the largest systems are now being designed, or re-routed, to avoid contested chokepoints and rival-controlled waters.

Taken together, the paper underlines that digital globalization generates new forms of strategic dependence. Because international data flows rely on geographically concentrated infrastructure controlled by a limited set of actors, economic integration can become a source of vulnerability.

Related literature. The paper relates to several active strands of research. First, by studying digital chokepoints and the potential weaponization of critical infrastructure, we contribute to the rapidly growing literature on geoeconomics (for surveys see Clayton et al., 2025b; Mohr and Trebesch, 2025). Existing work studies how dominant powers exert leverage through trade and financial interdependence (Clayton et al., 2026, 2025a; Broner et al., 2025; Mayer et al., 2025) and analyzes the design and effects of sanctions (Felbermayr et al., 2020; Nigmatulina, 2022; Egorov et al., 2025; Itskhoki and Ribakova, 2024). We extend this literature to the vulnerability and sabotage of physical infrastructure, a tool of geoeconomic statecraft that has been increasingly used in recent years but has not been studied quantitatively in this line of research.[1] We also provide a 175-year historical perspective that places contemporary concerns about digital coercion in a much longer-run context.

Second, we contribute to the vibrant literature on the welfare gains generated by data flows and digital connectivity. A growing macro-finance literature treats data as a nonrival production input with important implications for productivity, firm size, market structure, and economic measurement (Jones and Tonetti, 2020; Farboodi and Veldkamp, 2020; Veldkamp and Chung, 2024; Begenau et al., 2018; Brynjolfsson et al., 2025; Bhutani et al., 2026). AI and machine learning have further increased the economic value of data and the bandwidth that moves it (Goldfarb and Tucker, 2019). With a view on globalization, Hjort and Poulsen (2019) show that the arrival of subsea internet cables in Africa increased employment, firm entry, and incomes. Steinwender (2018) and Juhász and Steinwender (2018) study the rollout and economic effects of the nineteenth-century telegraph

[1] See e.g. The Economist, July 11, 2024: "How China and Russia could hobble the internet — The undersea cables that connect the world are becoming military targets."

network. This literature shares the view that digital connectivity has been a major source of welfare gains in recent decades. In this paper, we take the reverse view, focusing on the welfare losses and systemic risks that arise when the infrastructure underlying global data flows is disrupted. To our knowledge, this perspective is new to the economics literature.

Third, we contribute to a small but growing literature studying the subsea cable industry itself, including Caoui and Steck (2025) and Jeon and Rysman (2025). More recently, Imbs and Pauwels (2026) characterize the topology and fragility of the modern fiber-optic network, and Ottaviano (2026) uses cable-network centrality measures to study how digital infrastructure shapes the distance elasticity of goods trade. See Hjort (2025) for a review of the economics of the internet. Our paper differs from this literature by combining a long-run historical perspective with a focus on the geoeconomics of communication networks, including hegemonic dominance, cable ownership, and country-level exposure to disruption.

Finally, we contribute to the economic literature on networks and infrastructure. Existing work has primarily focused on overland transportation networks such as highways, railroads, and pipelines (e.g., Donaldson, 2018; Allen and Arkolakis, 2022; Fajgelbaum and Schaal, 2020; Gallea et al., 2022). Subsea cables and cross-border telecommunications infrastructure have received far less attention, despite their central role in globalization. We adapt traffic-assignment models from transportation economics to the global cable network. This allows us to trace data flows across individual cables and to quantify country-level dependence and vulnerability.

The remainder of the paper is structured as follows. Section 2 provides historical and technical background on undersea telecom cables, including a history of cable espionage and digital blockades via cable cutting, sanctions and embargoes. Section 3 introduces the *GlobeTel* dataset. Section **??** develops the network model and defines our welfare-based measures of vulnerability and dependence. Section 6 presents the main findings. Section 7 concludes.

2 Submarine cables: background and history

In the 21st century, near-instant global communication is taken for granted. It enables core aspects of the modern economy such as global value chains, e-commerce, high-frequency trading, and real-time teleconferencing. For most of history, however, communication was constrained by physical travel. Messages moved only as fast as horses, trains, or ships. Cross-border communication took days or even weeks. This began to change in 1850, with the first international submarine telegraph cable between England and France. Since then, the world has been transformed by a growing infrastructure of undersea cables and, more recently, satellites. Today's internet depends on around 500 submarine cables and several thousand satellites, built through both private and public investment.

In this section we briefly summarize the history of international telecommunication, with a particular focus on subsea cables and their (geo-)political disruptions and sabotage.

2.1 A primer on telegraph, telephone and internet cables, as well as satellites

The first permanent transatlantic cable started operations on July 27, 1866, after several prior, short-lived, and ultimately failed attempts. Over the next 40 years, telegraphy investments boomed and all major countries were connected to the growing global network, enabling same-day communication across much of the globe.

Since then, cable technology and bandwidth have advanced dramatically. Telegraph cables transmitted electrical pulses encoding Morse code, initially allowing only one message at a time and requiring manual decoding. Telephone cables allowed for analog voice transmission, later upgraded to digital signals, so that many simultaneous calls became possible, albeit with still limited capacity. In contrast to these earlier copper cables, today's fiber-optic cables use light to transmit massive volumes of data at near light speed, supporting millions of simultaneous connections. More specifically, in the 1890s, a typical transatlantic cable could transmit about 40 words per minute; by the 1950s, a typical cable supported 40 simultaneous phone calls, and 40,000 by the 1990s. Today's fiber-optic cables can transmit millions of high-quality video calls simultaneously. Remarkably, however, the basic principles of subsea cables have remained unchanged for over 150 years: cables are laid on the ocean floor, typically with a diameter similar to that of a garden hose (for a review of the history of telegraph cables see Ahvenainen (1981, 1996, 2004, 2011)).

Satellites have also seen a remarkable technical evolution since the launch of Sputnik in 1957. Early models had limited bandwidth and short lifespans, used mainly for military surveillance as well as broadcasting. Today, thousands of active satellites provide global coverage for communication, navigation, and earth observation. A notable recent development is the rise of large low-earth orbit (LEO) constellations like SpaceX's Starlink, which offer lower latency (lower transmission delay) and higher bandwidth compared to traditional geostationary satellites — those that orbit at fixed positions relative to the face of the earth.

When comparing cables and satellites it is clear that subsea cables remain the dominant, indispensable backbone of international data transmission. The key limiting factor of satellites is bandwidth, with additional concerns about their sensitivity to weather conditions, high latency, electronic jamming, and high costs.

It is striking how much more powerful subsea cables are compared to satellites, despite satellites receiving far more attention in the media and public debate. The most advanced fiber-optic subsea systems, such as Google's transatlantic *Grace Hopper* or Facebook's 45.000 km long *2Africa* cable are each designed to carry 500–800 terabits (tbts) per second. This far exceeds the combined bandwidth capacity of all satellites currently in orbit. For example, Starlink's constellation of over 7000 satellites currently has an estimated combined capacity of about 350 tbts per second.[2] In addition, satellite lifespans typically average around 5 years, compared to 25 years or more for subsea cables.

These technical constraints and higher costs make satellites a complement rather than a substitute for submarine cables. One of the main advantage of satellites is their relatively good coverage in remote regions with no or limited cable connectivity (such as rural regions, or in the middle of the ocean), as well as in conflict regions such as Ukraine.

Compared to terrestrial fiber-optic cables, the building of undersea cables is also typically much more cost-effective in the case of international connections. Undersea cables avoid the often complex and time-consuming process of securing permits, navigating border crossings, and constructing infrastructure across densely populated areas. In contrast, subsea cables can often be laid along direct, uninterrupted paths across the seabed, thus bypassing the regulatory and physical obstacles that complicate overland construction.

While technological advances have increased submarine cable capacity, rising demand driven by cloud computing, generative AI, and video streaming has also significantly increased reliance on that capacity. As a result, the physical geography of submarine cables continues to constrain global data flows, and the disruption of even a single cable can have major consequences. In 2013, a severed

[2] See Starlink's 2024 progress report `https://www.starlink.com/public-files/starlinkProgressReport_2024.pdf`.

cable near Alexandria caused a 60% drop in Egypt's internet speeds. In 2008, damage to key cables near Egypt disrupted services across the Middle East and South Asia, affecting up to 75 million users in India alone. In 2020, simultaneous faults in the WACS and SAT-3 cables caused widespread outages across West and Southern Africa, affecting tens of millions of people. More recently, in 2024, Houthi attacks in the Red Sea resulted in the severing of three subsea cables, disrupting access across East Africa and impacting an estimated 25% of European and South Asian internet traffic. These incidents reveal how submarine cables remain vulnerable chokepoints in the global internet infrastructure.

2.2 Digital blockades: evidence from 150 years

The physical characteristics of subsea cables make them vulnerable to accidental damage, but also to sanctions, sabotage, and espionage. This subsection provides a brief overview of "digital blockades" and espionage linked to geopolitical tensions, including periods of war. To compile this event archive, we build on our dataset as well as cable-specific documentation in our original sources such as the cable nomenclatures. Important general sources on these events include Kunert (1962), Headrick (1981, 1991b), Sontag and Drew (1999), Winkler (2008), Kennedy (2014), Starosielski (2015), Giles (2016), Bruton (2017).

A comprehensive account of these episodes, together with the underlying sources, is provided in Appendix A. It is organized by the three channels through which control over cables has been exercised: sanctions and the exclusion of adversaries (Section A.1), espionage in peacetime and war (Section A.2), and cable cutting and sabotage (Section A.3). The appendix spans the full 150-year history of the international cable system, from the American Civil War to the Baltic Sea and Taiwan incidents of the 2020s.

Two main takeaways emerge from our case review. First, throughout history, subsea cables have been repeatedly "weaponized" for geopolitical purposes — through sanctions, wire tapping, or cable cutting. Second, such actions have predominantly been carried out by great powers, in particular the UK and the US, but also China, Germany, and Russia.

3 The GlobeTel dataset — mapping the global telecom network 1850-2025

This section introduces the *GlobeTel* database on global telecommunication networks since 1850. It is the first database reconstructing a systematic history of every known international cable ever laid under the sea, and every telecommunication satellite ever launched into space. *GlobeTel* considerably expands on existing datasets with information from several primary sources which were digitized and coded by our research team in the span of more than three years. Figure 1 gives an overview of the sources used in constructing the database. The result is a database with detailed information on the technological, geographic, and economic characteristics of these critical pieces of infrastructure, providing a granular perspective on the historical evolution of one of the most extraordinary achievements of mankind — the global telecommunication network.

The database is the result of several years of manual data collection and synthesis. The full database currently covers 1,773 cable systems in history and today, with a total length of 2.7 million kilometers, 544 owning entities, and unique 2,935 landing points, as well as 21,234 satellites. Note that cable systems typically consist of several individual cables, i.e. cable segments. The earliest cable in the database is recorded in 1850, with the latest becoming active in 2025. Of the cable systems, 1,054 are active in the telegraph era (before 1950), 349 are active in the telephone era (1951 - 1994), and 370 are active in the internet era (1995 - 2025).

3.1 Subsea Cables (Telegraph, Telephone, Internet), 1850-2025

In the cables section of *GlobeTel*, we aim to cover the history of every international submarine telecommunication cable ever built. For each cable, the database records the dates of its entry into and exit from service, its geography (i.e. the places that it connected), which firms were involved in its construction, operation, and maintenance, the distribution of its ownership, the construction costs, the cable's capacity, and other technical features. Figure 2 gives an overview of the total active extent of the cables network across time.

The database is constructed manually from the underlying sources. We first coded each source independently, before combining them into a consensus database, containing the most authoritative version of each cable. The choice between following different sources was made case by case, weighing factors like

Figure 1. GlobeTel: data sources and coverage

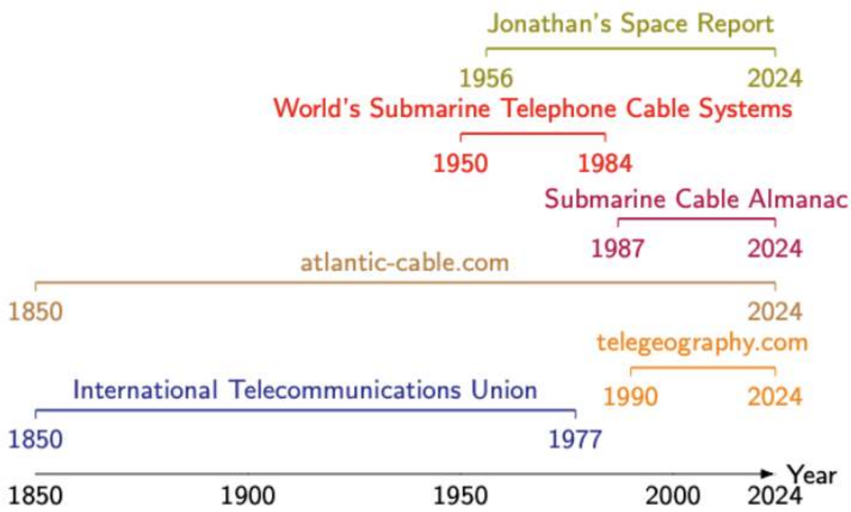

Figure 2. Total length of active cables per year

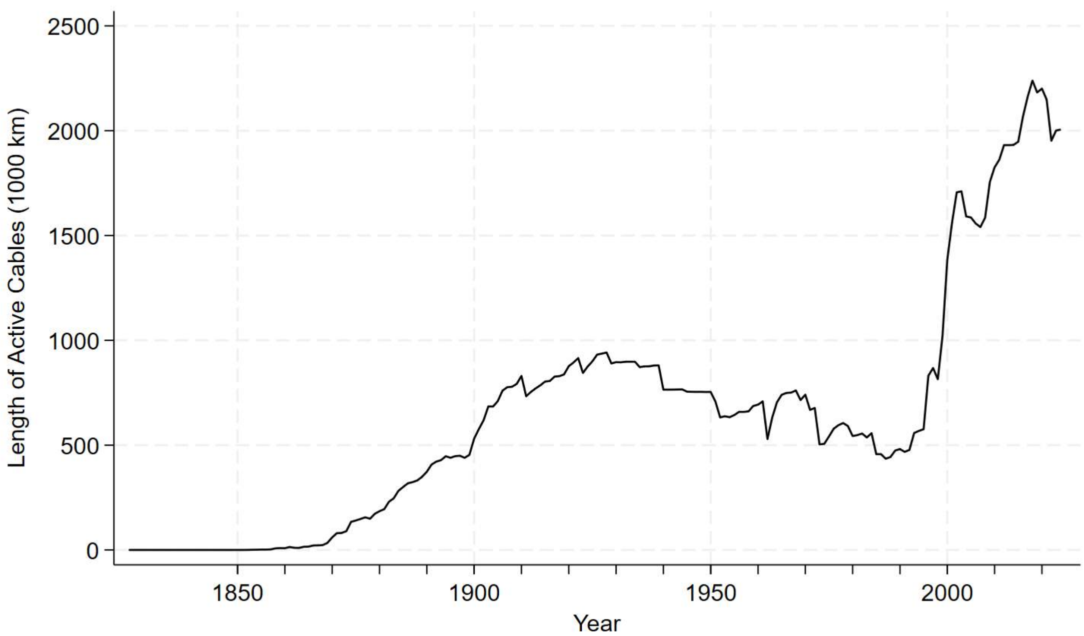

kind of source (official, commercial, personal), date of source (contemporary or later), and other factors.

The database has limitations in some dimensions, mostly due to lack of availability of source material. It was particularly difficult to find systematic cable-level data on exact ownership shares of different owner entities per cable system[3], costs of building[4], out-of-service dates[5], and bandwidth[6].

The cables database consists of five datasets:

- *Cables*: time-invariant features of cables
- *Landing points*: geographic information about every landing point of every cable
- *Capacities*: information on the capacities of each cable over time
- *Companies*: information on every company or government that was ever involved with a cable (owners, suppliers, operators, etc.)
- *Ownerships*: linking owners to cables across time

Data Sources. Given the wide time coverage of the database on cables, we must rely on several sources, some produced by official bodies, some by research companies, and others by volunteers and enthusiasts. The combination of all of these sources enables us to produce the most comprehensive and detailed database on international submarine cables. The only existing source which spans from the first telegraphy cables to the latest internet cables is *Atlantic Cable*, a freely accessible website. However, because the historical coverage of *Atlantic Cable* is limited, we also include information from the *Nomenclature des*

[3]Ownership shares that are different from 0, 50, or 100% are a relatively recent phenomenon, emerging mostly since around the 1990's, in the era of large-scale inter-continental cable systems frequently operated by groups of owners. In cases of unobserved ownership shares, we currently assume equal ownership of the unknown ownership portion between known owners.

[4]This variable can be approximated reasonably well, but precise cable-level costs of building are not always available.

[5]Ready-for-service dates are frequently observed and explicitly recorded for each cable, while out-of-service dates can sometimes only be bounded, e.g. in cases where a given cable is known to have gone out of service between the publishing of two subsequent censuses. In the absence of more specific cable-level information, we assume that a cable stays active until 25 years after the last time we observe that it is active.

[6]Especially for modern fiber-optic cables, the difference between potential capacity (i.e. the maximum theoretical capacity) and lit capacity (the capacity enabled by the hardware currently installed on either side of the cable) is important. Typically, information on potential capacity is more readily available than information on lit capacity, since lit capacity can frequently change for the same cable over time. The same distinction can be made for telegraph cables, albeit this is less of an issue.

câbles formant le réseau sous-marin du globe These documents are the results of a series of censuses of submarine cables performed (roughly) every five years between 1877 and 1977 by the International Telecommunication Union (ITU). Additionally, we add information from the *World's Submarine Telephone Cable Systems*, a source curated by the US Department of Commerce between 1950 and 1990. We further use additional information from a proprietary dataset on submarine cables maintained by *Telegeography*, one of the leading research companies on international telecommunications.[7] Our final source is the *Submarine Cables Almanac* by *SubTel Forum*, an ongoing effort started in 2011 with the aim to perform quarterly censuses of every active international submarine cable.

3.2 Satellites, 1957-2025

The satellites part of *GlobeTel* covers the history of every satellite that was ever used for commercial telecommunication. For each satellite, the database records the dates of its entry into and exit from service, details on its orbit, which firms and government agencies were involved in its construction, launch, operation, and maintenance, the distribution of its ownership, and other technical features. As future steps, we will also add information on every satellite's footprint (the geographic region covered by its signal), the construction and launch costs, and the satellite's capacity.

Similar to the cables part of the database, the satellites part of the database is also structured into five datasets:

- *Satellites*: time-invariant features of satellites
- *Orbits*: orbital information that can be used to calculate the satellite's geographic coverage
- *Capacities*: information on the capacities of each satellite across time
- *Companies*: information on every company or government that was ever involved with a satellite (owners, suppliers, space agencies, etc.)
- *Ownerships*: linking owners to satellites across time

Data Sources. Our main data source on satellites, covering our entire sample period, is the *General Catalog of Artificial Space Objects (GCAT)* by Jonathan C. McDowell. This dataset contains information on the technical aspects of satellites,

[7] Telegeography's EULA does not allow redistribution of the underlying data.

their out-of-service dates and ownership changes. We use two complementary data sources. First, the *Table of Artificial Satellites Launched*, periodically published by the ITU between 1967 and 1993. This source includes detailed information on every satellite launched before 1993. Second, the Czech Academy of Sciences curates a website with detailed technical data about every satellite that was launched into space between 1957 and 2012, including information about their bandwidth and capacity.

Figure 3. Number of satellite launches by year

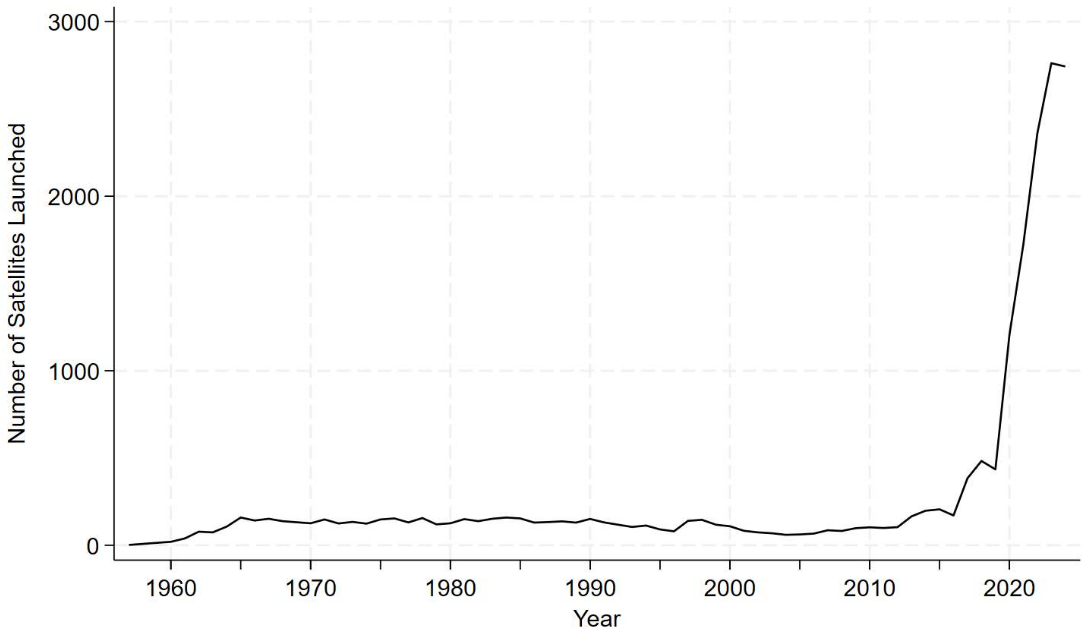

4 A model of data flows, dependence, and vulnerability

This section develops a simple model to quantify how exposed each country is to disruption of, and coercion through, the global telecommunications network. Its premise is that firms use the network to access productivity-enhancing knowledge that is dispersed geographically. Section 4.1 sets out the firms' demand for that knowledge, taking the cost of reaching each source as given. The next three subsections endogenize that cost: Section 4.2 describes the network and the traffic a routing of the data generates; Section 4.3 adds congestion and competitive carriers, which turn traffic into the cost of access; and Section 4.4 closes the model with an equilibrium tying demand, routing, and congestion together.

Section 4.5 then defines a country's vulnerability and dependence as the welfare cost of two network restrictions — a cable cut and an embargo — and pairs each with a simpler benchmark indicator.

4.1 Data flows

Each country n has a representative firm that produces with a Cobb-Douglas technology combining capital, production labour, and a bundle of knowledge inputs I_n,

$$Y_n = I_n^{\eta} F(K_n, L_n^p)^{1-\eta}, \qquad \eta \in (0,1), \tag{4.1}$$

where F is a constant-returns function of capital and labour and η is the elasticity of output with respect to the knowledge input.

Knowledge is differentiated by origin. Let x_{nj} be the data n draws from source j. The knowledge bundle is a CES aggregate of such data flows

$$I_n = \left[\sum_j B_j^{1/\rho} x_{nj}^{(\rho-1)/\rho}\right]^{\rho/(\rho-1)}, \qquad \rho > 1, \tag{4.2}$$

where ρ is the elasticity of substitution across sources and $B_j > 0$ is the productivity of j's knowledge: the higher it is, the more output a unit of data sourced from j yields.

Acquiring data from a source is costly to the firm in two ways. First, it pays a telecommunications carrier a price p_{nj} for each unit of data carried from j. Second, acquiring data takes time: workers are diverted from production as they are held up by loading websites, lagging calls, etc. We capture this by assuming that receiving one unit of data from j absorbs $\chi\, \tau_{nj}$ worker-hours, where τ_{nj}, the connection's *latency*, is the round-trip time of an exchange with j: the firm's query travels out to j, and the requested data travel back. A higher-latency (slower) link ties up more worker time per unit of data. Acquiring x_{nj} units therefore diverts $\chi\, \tau_{nj}\, x_{nj}$ worker-hours from production. Thus, aggregating across sources, knowledge acquisition claims labour $L_n^I = \sum_j \chi\, \tau_{nj}\, x_{nj}$.

The two costs combine in the firm's budget. The firm hires capital K_n at the rental rate r_n and labour L_n at the wage w_n, and pays carriers for the data it draws, so its total cost is $r_n K_n + w_n L_n + \sum_j p_{nj} x_{nj}$. Labour serves either production or knowledge acquisition, $L_n = L_n^p + L_n^I$. Substituting $L_n^I = \sum_j \chi\, \tau_{nj}\, x_{nj}$ and collecting the terms that multiply each x_{nj}, total cost can be rewritten as

$$r_n K_n + w_n L_n^p + \sum_j \left(p_{nj} + w_n \chi\, \tau_{nj}\right) x_{nj}.$$

The coefficient in parenthesis is what one unit of data from j adds to the firm's costs: the price paid to the carrier plus the wage bill of the worker time that unit ties up. We call it the *generalized cost of access*,

$$\kappa_{nj} = p_{nj} + w_n \chi \, \tau_{nj}, \tag{4.3}$$

the price the firm effectively pays for a unit of knowledge from j. Note that neither the carrier's price p_{nj} nor the latency τ_{nj} is a primitive of the model: we derive them from optimal routing and network topology in Sections 4.2–4.4.

The firm assembles the knowledge bundle at least generalized cost. Cost-minimizing (4.2) gives the demand for data from each source,

$$x_{nj} = I_n \, B_j \, \kappa_{nj}^{-\rho} / \left(P_n^I\right)^{-\rho}, \qquad P_n^I = \left[\textstyle\sum_j B_j \, \kappa_{nj}^{1-\rho}\right]^{1/(1-\rho)}, \tag{4.4}$$

where P_n^I is the CES price of the knowledge bundle faced by country n. Under Cobb-Douglas (4.1), knowledge outlays are the share η of revenue, $P_n^I I_n = \eta Y_n$, so

$$x_{nj} = \eta \, Y_n \, B_j \, \kappa_{nj}^{-\rho} \, \left(P_n^I\right)^{\rho-1}. \tag{4.5}$$

Equation (4.5) is a gravity equation for data flows: the data n draws from j rises with the productivity B_j of j's knowledge and with n's size Y_n, and falls with the generalized cost of access κ_{nj}.

4.2 Routing and traffic

The network is a graph of nodes joined by links. A link ℓ is a connection (submarine, terrestrial, or wireless) between two nodes. A path p on such network is a round trip between n and j: a sequence of links carrying the firm's query out to j and the data back to n. We let $\mathcal{P}(n, j)$ denote the collection of all such paths. Telecom carriers route data flows on this network, where routing the flow x_{nj} means that such flow is assigned to a path. We let $\pi_{nj}(p) \geq 0$ denote the share of x_{nj} carried on path p. A routing induces a link share $\pi(\ell, n, j) = \sum_{p \in \mathcal{P}(n,j):\, \ell \in p} \pi_{nj}(p)$, the fraction of x_{nj} that crosses link ℓ; summing over all flows, the *traffic* on a link is the total data it carries,

$$\text{Traffic}_\ell = \sum_n \sum_j \pi(\ell, n, j) \, x_{nj}. \tag{4.6}$$

We also write $\text{Traffic}_\ell^n = \sum_j \pi(\ell, n, j)\, x_{nj}$ for the *domestic traffic* of country n, the part of Traffic_ℓ linked to the knowledge n consumes.

Figure 4 illustrates a routing and the link traffic it generates. In the two-node network on the left, the single A–B flow loads every link it crosses equally; adding a node C on the right concentrates traffic on the links nearer A, which now carry A's flows to both destinations.

Figure 4. Graphical representation of Traffic Assignment

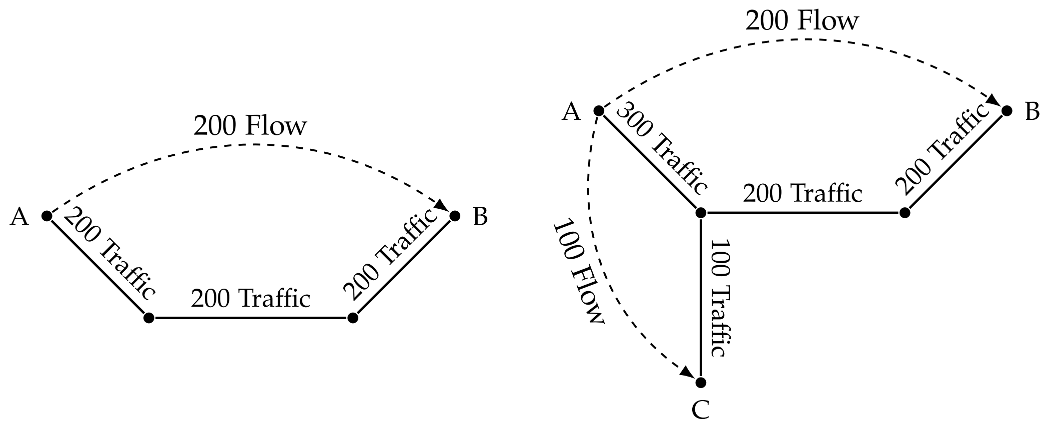

4.3 Congestion and carrier costs

Every time data cross a given link ℓ, the communication accrues a delay of t_ℓ; absent congestion this equals the link's *free-flow latency* t_ℓ^0. However, we assume that latency rises with the traffic the link carries. Specifically, we let Cap_ℓ denote the capacity of link ℓ and we assume that latency increases with Traffic_ℓ, and especially so when Traffic_ℓ exceeds Cap_ℓ. We capture this with the standard congestion function

$$t_\ell(\text{Traffic}_\ell) = t_\ell^0 \left[1 + \alpha \left(\text{Traffic}_\ell / \text{Cap}_\ell\right)^\beta\right], \tag{4.7}$$

with $\alpha \geq 0$, $\beta \geq 1$.[8]

We assume that a carrier's cost of moving a unit of data across link ℓ is $\varphi\, t_\ell$, where φ is the cost per unit of transit time. This captures the idea that every extra moment in the network ties up energy, equipment, and transmission capacity, and thus increases costs. Any path p has a route latency $\tau_{nj}(p) = \sum_{\ell \in p} t_\ell$, the

[8]Equation (4.7) is the Bureau of Public Roads function; α and β are estimated from within-day congestion rhythms (Section 5; Appendix Appendix C).

latency accumulated over the round trip; so, carrying a unit of data along path p costs the carrier $\varphi\,\tau_{nj}(p)$.

We assume the telecom market is perfectly competitive, which pins down both prices and routes. Carriers make zero profit, so each path is priced at cost; by (4.3), the generalized cost of sourcing a unit from j along a path p is thus

$$\kappa_{nj}(p) = (\varphi + w_n\chi)\,\tau_{nj}(p) \equiv c_n\,\tau_{nj}(p), \tag{4.8}$$

with $c_n \equiv \varphi + w_n\chi$ a positive number specific to country n. Moreover, perfect competition implies that each carrier takes the link traffic, and hence the latencies, as given and offers firms the path that induces the lowest generalized cost of access, i.e.,

$$\kappa_{nj} = \min_{p\in\mathcal{P}(n,j)} \; c_n \sum_{\ell\in p} t_\ell(\text{Traffic}_\ell), \tag{4.9}$$

where $t_\ell(\cdot)$ is the congestion function (4.7); we write τ_{nj} for the latency of the least-cost path. The link traffic, and hence which route attains the minimum, is pinned down by the equilibrium of Section 4.4.

4.4 Equilibrium

Through congestion (4.7), the cost of a route depends on the routes every other flow takes, so demands, routing, and latencies must be determined together. An *equilibrium* is a set of data flows $\{x_{nj}\}$, a routing $\{\pi_{nj}(p)\}$, and link latencies $\{t_\ell\}$ such that:

1. *Optimal knowledge sourcing.* taking as given the equilibrium generalized costs $\{\kappa_{nj}\}$, each firm assembles its knowledge bundle at least cost, so the flows satisfy the gravity equation (4.5);

2. *Competitive routing.* taking the link traffic as given, perfectly competitive carriers price every path at cost and carry each flow along a path of least generalized cost, so the routing attains the minimum in (4.9);

3. *Market clearing.* for each (n, j) every flow is delivered in full,

$$\sum_{p\in\mathcal{P}(n,j)} \pi_{nj}(p) = 1, \tag{4.10}$$

   and the link latencies are consistent with the congestion function (4.7) evaluated at the equilibrium distribution of traffic (4.6).

For given demand, conditions (ii)–(iii) yield what the transport literature calls a *user* (or Wardrop) *equilibrium*. The equilibrium is a fixed point: latencies shape demand and routing, demand and routing determine link traffic, and traffic feeds back into latencies through congestion.

Our equilibrium sits within a large literature on routing flows over economic networks. One strand assigns flows to least-cost paths whose link costs are fixed (Donaldson and Hornbeck, 2016; Allen and Arkolakis, 2014; Arkolakis and Walsh, 2023). Our model nests this case: with $\alpha = 0$, costs are independent of traffic. A second, more recent strand keeps the congestion feedback, as we do, but closes routing differently. Fajgelbaum and Schaal (2020) take a planner's perspective. Allen and Arkolakis (2022) model routing as the result of cost minimization with agents who have route-specific preference shocks. Motivated by the competitive and decentralized nature of telecom networks, we instead consider a user equilibrium where firms (and thus carriers) do not care about the route their data takes but only about their generalized cost.

4.5 Measuring dependence and vulnerability

We define vulnerability and dependence as the welfare cost of two specific restrictions, each computed by perturbing the network and re-solving the equilibrium of Section 4.4:

- *Vulnerability* is the welfare cost of a cable cut. To measure country n's vulnerability to a set of cables, remove them, re-solve for the new equilibrium latencies, and measure the resulting welfare difference.
- *Dependence* of n on s is the welfare cost of an embargo by s: the change in n's welfare between the global routing equilibrium and the equilibrium in which n's data cannot flow on s-owned links. Deny n those links, re-solve, and measure the resulting welfare difference.

Both restrictions remove links from the network, forcing a re-routing of some data flows. The re-routing raises latency twice over: re-routed data takes longer paths, and the added congestion slows all other flows too, even those not re-routed. Both effects raise the welfare cost through the equilibrium latencies τ_{nj}.

Because we do not model the consumer side, we take a reduced-form approach and identify a country's welfare W_n with its real income. A restriction of the network (e.g., a cable cut or an embargo) raises the price of knowledge P_n^I.

By the envelope theorem, the first-order effect on W_n is

$$\Delta \ln W_n = -\,\eta\,\Delta \ln P_n^I \equiv \eta\, Z_n, \qquad Z_n \equiv -\Delta \ln P_n^I, \tag{4.11}$$

the elasticity η times the deterioration in access. Because the generalized cost is proportional to latency, $\kappa_{nj} = c_n \tau_{nj}$, the price index factors as $P_n^I = c_n \widetilde{P}_n^I$ with $\widetilde{P}_n^I \equiv \left[\sum_j B_j \tau_{nj}^{1-\rho}\right]^{1/(1-\rho)}$; a restriction moves latencies but not c_n, which cancels:

$$Z_n = -\,\Delta \ln \widetilde{P}_n^I = -\,\frac{1}{1-\rho}\,\Delta \ln \left[\textstyle\sum_j B_j\, \tau_{nj}^{1-\rho}\right]. \tag{4.12}$$

Sufficient statistics. Equation (4.12) appears to require the knowledge productivities $\{B_j\}$; in fact it does not. By the demand system (4.4) and the factorization $\kappa_{nj} = c_n \tau_{nj}$, the share of source j in country n's knowledge outlays is $s_{nj} = \kappa_{nj} x_{nj} / \sum_k \kappa_{nk} x_{nk} = \tau_{nj} x_{nj} / \sum_k \tau_{nk} x_{nk}$, a latency-weighted flow share observed directly in the baseline equilibrium. Writing $\widehat{\tau}_{nj} = \tau'_{nj}/\tau_{nj}$ for the proportional change in latency the restriction induces, (4.12) becomes

$$Z_n = \frac{1}{\rho-1}\,\ln \left[\textstyle\sum_j s_{nj}\, \widehat{\tau}_{nj}^{1-\rho}\right], \tag{4.13}$$

so the welfare measure requires only the baseline shares, the counterfactual latency changes, and the elasticities η and ρ: neither the productivities $\{B_j\}$ nor the cost scalars φ, w_n, χ enter, and the counterfactual equilibrium itself can be solved in relative changes from the observed flows, in the tradition of exact hat algebra (Dekle et al., 2007). For small restrictions, (4.13) reduces to $Z_n \approx -\sum_j s_{nj}\,\Delta \ln \tau_{nj}$: the welfare cost is η times the share-weighted increase in latency, so ρ matters only at second order, once substitution across sources starts to mitigate large shocks — the headline numbers are not fragile to its exact value, and the benchmark indicator below is the routing-based counterpart of the same weighting. Finally, if a restriction disconnects n from j outright ($\widehat{\tau}_{nj} = \infty$), the j-term drops out of (4.13) and the loss stays finite at $\frac{1}{\rho-1} \ln \frac{1}{1-s_{nj}}$ log points of P_n^I: with $\rho > 1$ sources are substitutes and no single one is essential. Appendix Appendix C sets out, step by step, how to estimate these objects and turn them into the welfare measures.

Aggregating across countries. Equation (4.11) prices a restriction for one country in log points of its real income; summarizing a restriction's global cost requires saying how countries add up. Write $\Delta \ln P_n = -Z_n \geq 0$ for the increase in country n's price index under the restriction. We aggregate as a money metric: each country's loss is first converted into the share of real income it destroys, $1 - e^{-\eta \, \Delta \ln P_n}$, these shares are scaled by economy size, and the sum is expressed relative to world income,

$$\Delta W \;=\; \frac{\sum_n \mathrm{GDP}_n \left(1 - e^{-\eta \, \Delta \ln P_n}\right)}{\sum_m \mathrm{GDP}_m}. \tag{4.14}$$

The numerator alone is the world loss in dollars; dividing by world GDP only changes units, making ΔW a pure number on the same $[0, 1)$ scale as the country-level shares and comparable across years and scenarios.

Benchmark. We complement the welfare-based indicators with simpler benchmark measures of vulnerability and dependence. They rely on less of the model: they use only the equilibrium routing and traffic, not the demand parameters η, ρ, or the knowledge productivities $\{B_j\}$.

Holding routing fixed, the loss of one cable disrupts the data that crosses it. Country n's vulnerability to cable ℓ is thus the share of its total data flows ($\mathrm{TotFlows}_n = \sum_j x_{nj}$) that flows on ℓ; similarly, global vulnerability is defined as a share of global flows ($\mathrm{GlobalFlows} = \sum_{n,j} x_{nj}$)

$$\mathrm{Vuln}_n^\ell = \frac{\mathrm{Traffic}_\ell^n}{\mathrm{TotFlows}_n}, \qquad \mathrm{GlobVuln}^\ell = \frac{\mathrm{Traffic}_\ell}{\mathrm{GlobalFlows}}, \tag{4.15}$$

This benchmark is a short-run notion: it holds routing fixed, before carriers re-route around the lost cable. For the directly affected traffic it is therefore an upper bound on the cost, since re-routing would restore part of the connection (but with a higher latency). Instead, for the other flows, this measure is a lower bound as it ignores the extra congestion that re-routed flows would impose on the rest of the network.

Figure 5 illustrates the vulnerability indicator. We see that vulnerability is highest where traffic concentrates on few cables (low redundancy): corridor C1 carries all flows over only two cables, so losing either disrupts half the data crossing it.

While vulnerability measures n's exposure to the loss of a cable, dependence applies the same logic to *ownership*. Specifically, dependence measures a country's exposure to the set of cables owned by another country, which could be denied via sanctions or embargoes. We measure country n's dependence on country s

Figure 5. Graphical representation of Vulnerability

(a) $GlobalVulnerability^{C1} = \frac{1/2 \times 300}{300} = 50\%$

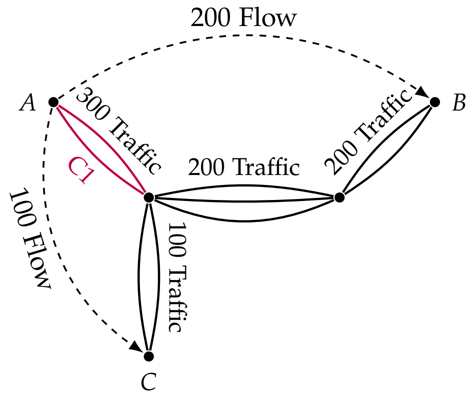

(b) $GlobalVulnerability^{C2} = \frac{1/3 \times 200}{300} \approx 22\%$

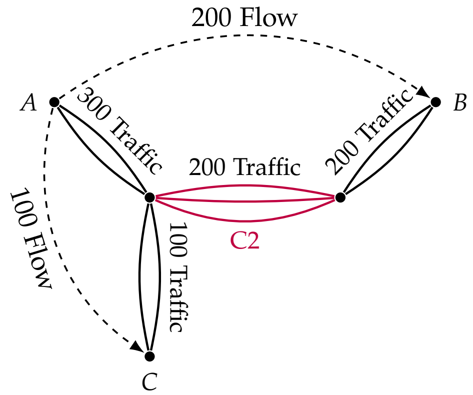

(c) $GlobalVulnerability^{C3} = \frac{1/2 \times 200}{300} \approx 33\%$

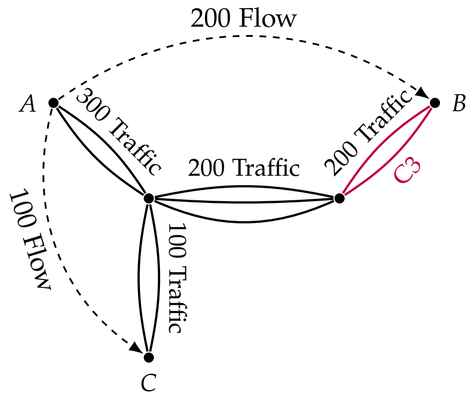

(d) $GlobalVulnerability^{C4} = \frac{1/2 \times 100}{300} \approx 17\%$

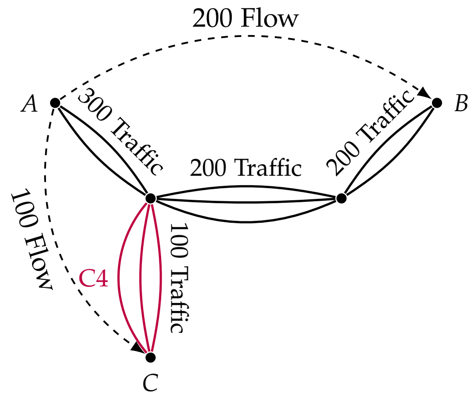

as the share of country n's data that crosses cables owned by s. Formally,

$$\text{Dep}_n^s = \sum_j \left(\max_\ell \ \pi(\ell, n, j) \cdot \mathbf{1}\{\ell \text{ owned by } s\} \right) \frac{x_{nj}}{\text{TotFlows}_n}, \tag{4.16}$$

while dependence of the global economy on country s is the share of global data that crosses cables owned by s,

$$\text{GlobDep}^s = \sum_n \sum_j \left(\max_\ell \ \pi(\ell, n, j) \cdot \mathbf{1}\{\ell \text{ owned by } s\} \right) \frac{x_{nj}}{\text{GlobalFlows}}. \tag{4.17}$$

As with vulnerability, these indicators are a short-run notion that holds routing fixed: faced with an embargo, n would re-route around s where it could, reducing its cost and effectiveness. The indicator also measures a country's exposure to *espionage*: a country s that carries n's data can read it as well as block it. Both readings are consistent with the max operator, as interception or garbling is non-rival across segments — once s has seen or blocked a packet, a second s-owned link on the same route adds nothing.

Figure 6 illustrates the measure: the blue links are owned by country B, and the panels vary that ownership; once B controls the link nearest it, a second B-owned link on the same route leaves dependence unchanged.

5 Taking the Model to the Data

Before we run our counterfactuals, we briefly discuss how we bring the model to the data, estimate its parameters, and validate its main predictions. A more detailed discussion is in the Appendix.

Observables: demand and capacity. The first step in bringing our model to the data is to consistently measure international bilateral data flows. For this we rely on the used international capacity (bandwidth demand) measure provided by TeleGeography.com, which captures the expected peak volume of data exchanged between two countries, net of transit by or to third parties.[9] To build the demand shares that enter our welfare analysis, we complement this measure with a country-level estimate of domestic traffic, defined as total backbone traffic net of

[9]The flow data are undirected, and we split pair totals equally between the two directions. Under the gravity equation (4.5) the split is exact when B_j is proportional to j's demand shifter — equivalently, when bilateral flows are balanced pair by pair; robustness replaces it with a split proportional to directed trade. To first order, the world aggregate (4.14) is invariant to the split: misdirected shares only reallocate a pair's loss between its two endpoints.

Figure 6. Graphical representation of Dependence

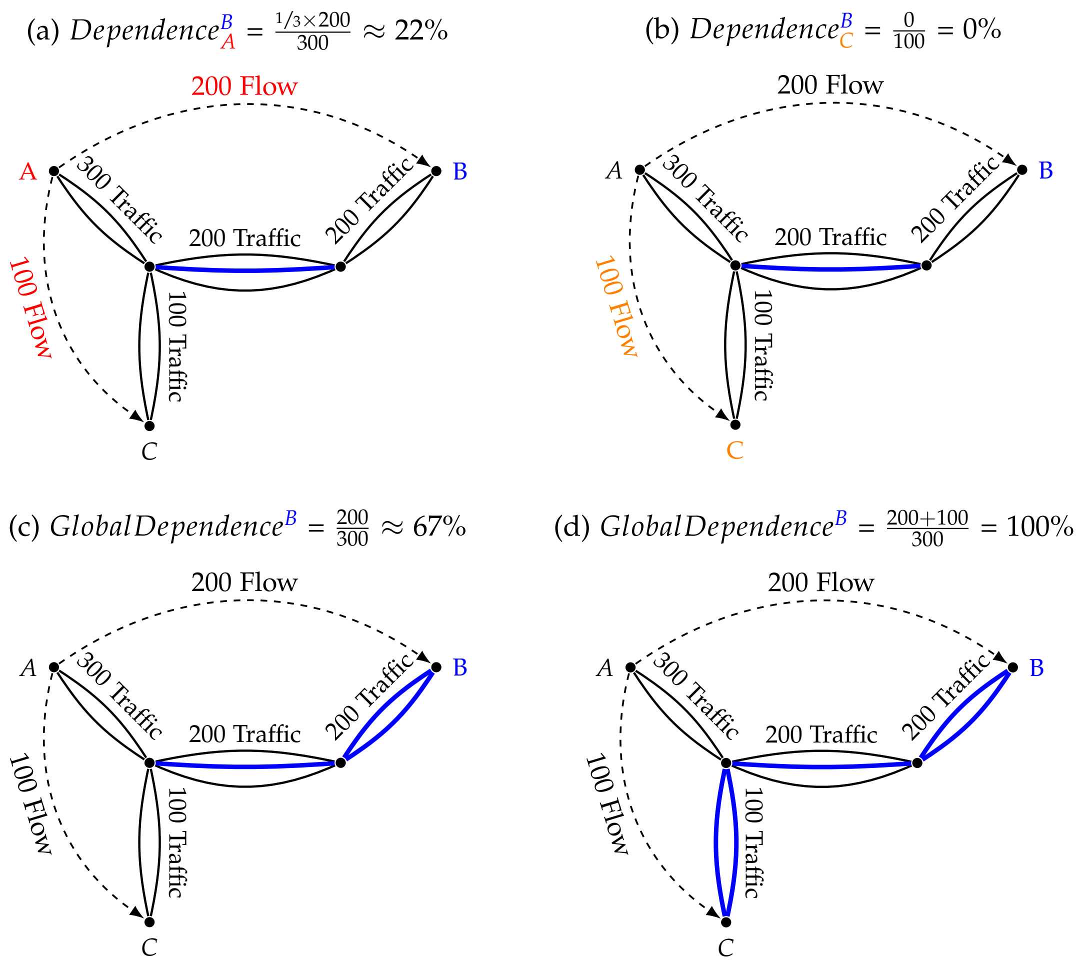

(a) $Dependence_{A}^{B} = \frac{1/3 \times 200}{300} \approx 22\%$

(b) $Dependence_{C}^{B} = \frac{0}{100} = 0\%$

(c) $GlobalDependence^{B} = \frac{200}{300} \approx 67\%$

(d) $GlobalDependence^{B} = \frac{200+100}{300} = 100\%$

estimated international traffic (also from TeleGeography.com). A second variable that enters the welfare calculations is the capacity of the network edges. For this we use the potential capacity measure of our dataset, which captures the theoretical transmission capability of a cable given its technical characteristics and the equipment available in a given year.

Observables: the physical network. We model submarine cables as items that allow data to be exchanged between any two of the locations they connect, along a pre-determined path. While the general topology of submarine cable systems is often public, the specific path they trace on the ocean floor is unobservable and often difficult to reconstruct. This limitation is important: an imprecise measurement of their path leads to imprecise measures of the effective distance between the landing points, an input to any network analysis. Based on the maps of various cable systems and the economics of cable laying, we assume that each system is a connected graph laid on the ocean floor and organized in tiered branches and landings (primary, secondary, tertiary, etc.). Specifically, a system's primary branches minimize the seafloor distance between a pair of primary landings, the secondary branches minimize the seafloor distance between a primary branch and a secondary landings, and so on. We assign tiers to landings and segments in order to maximize the fit between the predicted path and the observed topographic features of the system. A simple external check shows the fit is excellent: across the 604 systems of our model dataset, the median ratio of fitted to declared length is 1.03 (Appendix B.1).

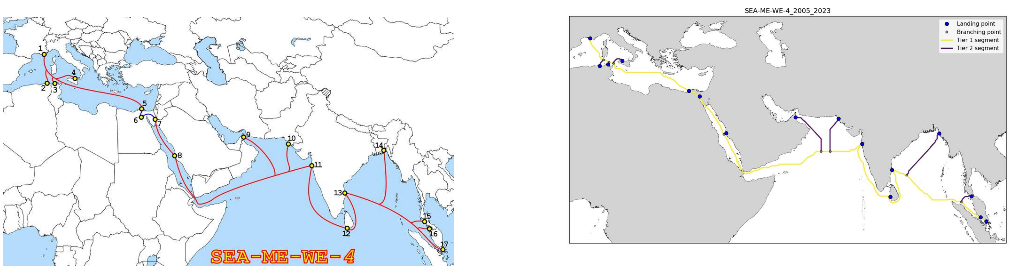

Figure 7. Example of a modern cable system, SEA-ME-WE 4. Left: map from the cable owner's website. Right: path as reconstructed by our algorithm.

A second challenge in our analysis is the lack of information on international terrestrial networks. While most countries have a fairly developed terrestrial telecommunications backbone, most of it is explicitly reserved for and targeted to the needs of domestic traffic; moreover, the extent to which these domestic networks exchange traffic with each other is unclear. We address this limitation in two steps. First, to ensure that every origin and destination of data flows is connected to the global network, we add domestic links that connect to the country's own landing points. These links are domestic in the sense that they can only carry data that is originated from or destined to the country they serve. Second, we asses the existence of potential terrestrial transit cables with a data-driven approach. We use existing information on terrestrial systems and internet exchange points to construct a menu of candidate terrestrial overlays: regional "configurations" ranging from purely national corridors to dense international meshes. For each of these configurations we then simulate our model globally and derive the predicted minimum latency between every city pair in the latency dataset from wonderproxy.com. We then select the terrestrial configuration that provides the best fit between our predictions and the minimum yearly observed latency. The selected overlay cuts the prediction error by a third relative to assuming a submarine-only network ($R^2 = 0.90$ across 18,329 country-pair-years). Appendix B.2 reports the criterion and the full choice set for every region. All of this combines into Figure 8, representing the global telecom network that underlies our analysis. Blue lines represent our reconstructed path for each submarine cable, black lines represent the fit-maximizing international transit terrestrial network, red lines show the domestic links connecting each landing point to its country's origins or destinations of data flows, whenever it can be reached by land.

Finally, we add satellites to our network. We use publicly available orbital information on satellites and satellite constellations to infer their average yearly regional coverage and their altitude. We model satellites as infrastructure that can relay a signal between any two points that are simultaneously located within their coverage area.

Estimation: routing parameters t^0, α, and β. With the network in place, we turn to the congestion function (4.7) of Section 4.3 that turns traffic into latency, $t_\ell = t^0_\ell \left[1 + \alpha(\text{Traffic}_\ell/\text{Cap}_\ell)^\beta\right]$. Its free-flow latency t^0_ℓ can be written as the ratio of the length of edge ℓ to the speed at which the signal propagates in the absence of congestion. We use latency data to verify the common belief that signals travel at 50% of the speed of light in fiber-optic cables, meaning that a signal traveling on an uncongested strand of fiber covers 150 km in 1 millisecond.

Concerning the shape parameters, α and β, we start by imposing the restriction that latency doubles at twice capacity, which ties $\beta = -\log_2 \alpha$ and leaves

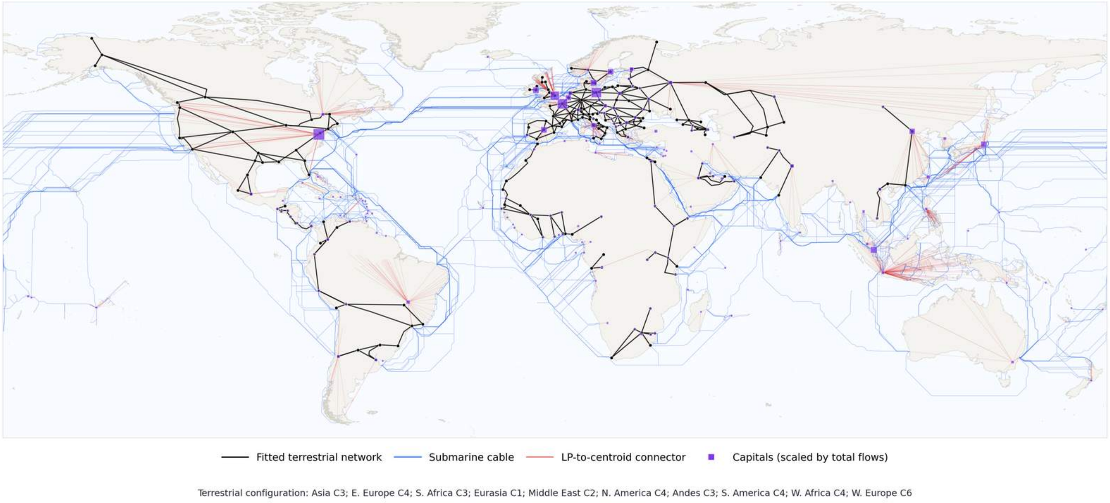

Figure 8. Global terrestrial network configuration that maximizes model's fit to latency data. Appendix B.2 shows the choice sets for all regions.

us with only one parameter to estimate, α. To estimate this parameter, which measures the impact of congestion on latency, we must address a potential endogeneity issue: heavily used routes are also the most optimized, implying that unobserved route quality would confound the relationship. To address this concern, we identify our parameters from intra-day changes in congestion that are linked to the daylight cycle. First, we download data from radar.cloudflare.com on the hourly traffic-to-peak ratio for 90 domestic networks. Second, we multiply that ratio by the international demand profile of that country. Third, we push these shifted profiles through the routing model to obtain a predicted hourly congestion rhythm for every origin–destination pair. If, within a pair and a day, route quality and equipment are constant, then the only reason observed latency should rise in peak hours is that capacity is being filled up (congestion). Solving the fixed point between the implied regression and the routing equilibrium gives $\hat{\alpha} \simeq 0.24$ and $\hat{\beta} \simeq 2.0$, stable across our sample period (Figure 9 and Appendix C.1).

Estimation: demand elasticities, ρ and η. The last parameters that enter our welfare analysis come from the demand side of the model. As is standard in the gravity literature, we estimate ρ—the elasticity of substitution across knowledge sources in the CES bundle (4.2)—from how bilateral data flows respond to the resistance term (latency), i.e. from the gravity equation (4.5). Of course, in our setting, observed latency is endogenous: traffic creates congestion, and latency

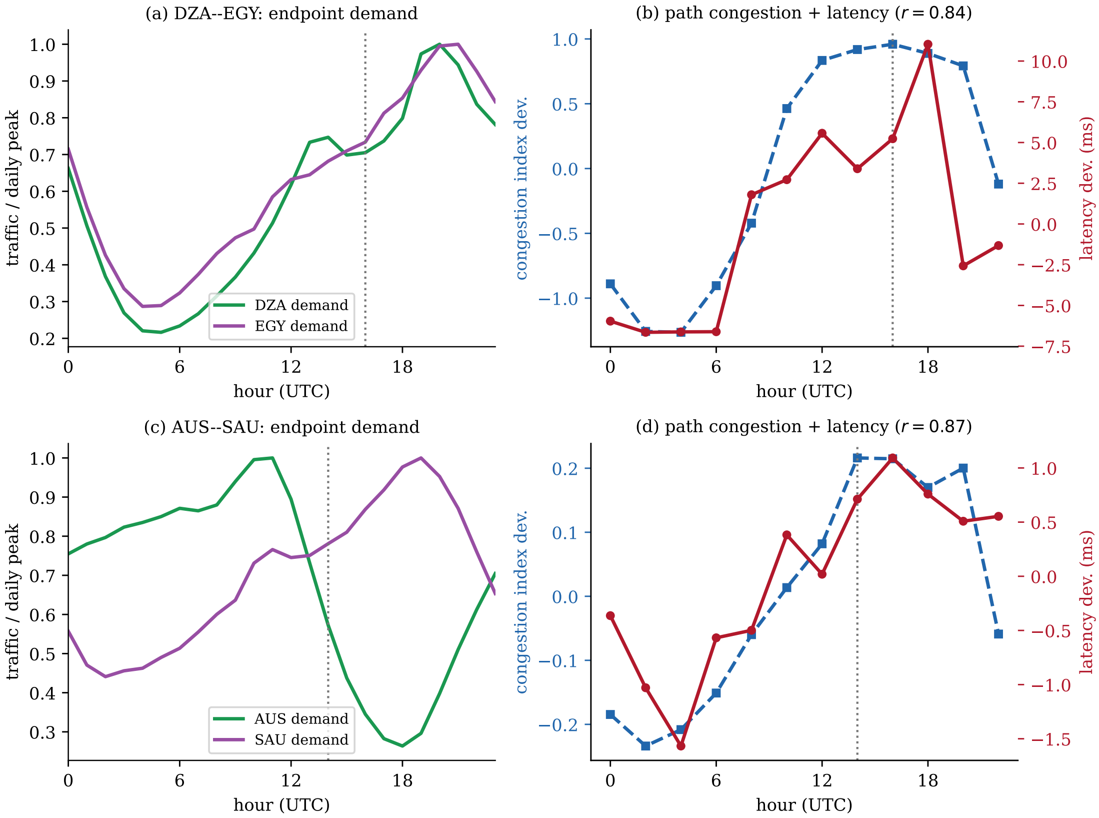

Figure 9. The traffic-latency rhythm. Left: the domestic traffic-to-peak profile of four countries. Right: correlation between the predicted hourly congestion along the route (accounting for third parties' transit) and the route's hourly excess latency.

deters traffic. Following the literature, we instrument observed latency with the free-flow (uncongested) latency of the physical network. Our preferred estimates put the elasticity between 2.5 and 3.5 (Appendix C.2).

For the other parameter—the output elasticity of data η in the production function (4.1)—a standard approach would be to calibrate it to a country's expenditure share in the telecommunications sector (around 1% of GDP). According to our model, however, this would only yield a lower bound, because it misses the shadow cost of the time workers spend acquiring data. As we show in Appendix C.3, the existence of countries whose domestic telecom market is shielded from international competition yields the identifying variation we need. Gate-keeping monopolists would optimally charge a markup when selling access to the international network, and this markup is a function of the elasticity of demand for data, η. Using data on the price of international city-pair-specific routing and domestic country-specific internet access, we estimate $\eta = 0.5$, inside a feasible set of [0.31, 0.59] pinned by a regulation-controlled floor and a demand-coherence ceiling (Appendix C.3).

Validation: testing the model's predictions. So far, we have estimated the model's parameters using data on latency, cable length, prices, and data flows. The assembled model makes sharp predictions we did not target in estimation, which we now confront with the data: the identity of cables that packets actually cross, the price of connectivity, and the capacity operators light.

For the former we use traceroutes, a diagnostic test that tracks the path of data packets across millions of routers scattered around the world. Existing traceroute datasets have incomplete paths and uncertainty in the geolocation of routers. Despite this noise, we find that traceroutes strongly validate the accuracy of our routing predictions. Comparing model-predicted and traceroute-implied cable usage shares yields a Spearman correlation of 0.64 across 172 cables. At the pair-cable level, a cable on the model's route for a city pair is observed for that pair 30.7 percent of the time, against 3.6 percent for cables the model does not select; among pairs with at least five traceroutes, pooled precision reaches 54.0 percent, against 14.3 percent under a geography-aware random benchmark that grants the same number of guesses from the same regional cable vocabulary (Figure 10). Finally, we use traceroutes to validate our model's finding that satellites are never used along the path except for the last mile. In fact, satellites are used in only 0.3% of the available traces, and 90% of that usage occurs on the first or last leg of the route they serve (Appendix C.4).

The model also prices data transport in proportion to latency—under com-

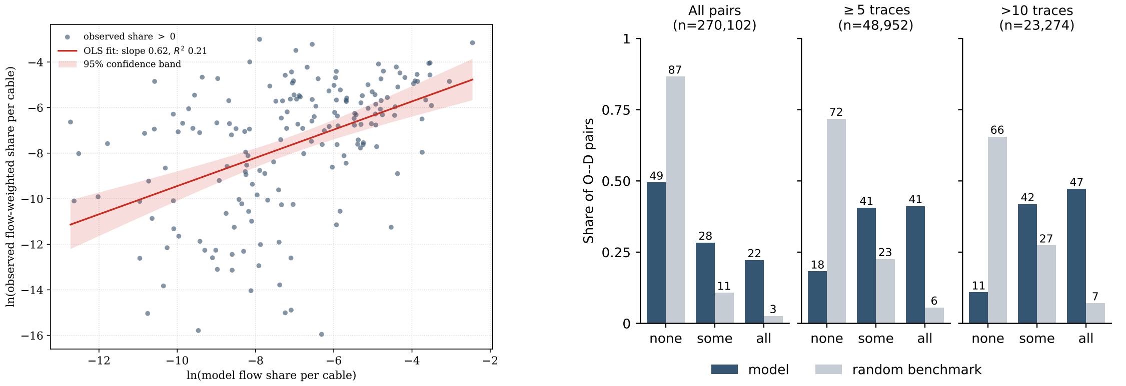

Figure 10. Traceroute validation of predicted cable use. Left: model-predicted versus traceroute-implied usage shares by cable. Right: the share of a pair's predicted cables confirmed by its traceroutes (none/some/all), for the model and for a geography-aware random benchmark. Appendix C.4 reports construction, denominators, and robustness.

petition a carrier charges $p_{nj} = \varphi\, \tau_{nj}$ (Section 4.3)—a shape restriction we can test directly. Leased-line prices bear it out: absorbing city fixed effects at both endpoints, so that access and termination costs drop out, the elasticity of price per Gbps with respect to latency is 1.00 (s.e. 0.08), and instrumenting realized latency with physical free-flow latency leaves the relationship intact—indicating that it reflects physical cost rather than congestion (Appendix C.3).

A final validation exercise for our routing predictions is based on lit-capacity data, i.e., the amount of a system's potential capacity that network operators keep "ready" for the transit demand they expect to receive. As long as this decision is responsive to expected volumes, we can use this variable to test whether it correlates with our model's predicted system-level transit volumes. Our regressions validate this conjecture and then show that our model can also accurately predict the utilization ratio of systems (lit over potential capacity), which is a much tighter test, as it cannot be driven by a spurious correlation with the presence of large-capacity systems (Appendix C.5).

6 Main Results

6.1 Hegemons dominate the global telecom network, giving them leverage in surveillance and sanctions

This section documents that hegemons have historically dominated the global telecommunication infrastructure. This pattern spans the entire history of long-distance communication, from the telegraph era to the present-day internet and satellite networks. In a first step, we document cable ownership across countries, showing the disproportionate role that the UK and the US have long held in this infrastructure, acting as backbone provides with truly a global network. Second, we show that ownership shares should be seen as a lower bound of dependence as the centrality of hegemonic cables implies that the majority of global data flows traverse its infrastructure at one point in their route. This gives hegemons considerable leverage for espionage and surveillance. Third, we simulate the welfare effects of cable sanctions, assuming that dominant cable owners like the United States bar rival countries from using their network. The results point to large welfare effects, especially in the great power's neighboring regions.

Ownership dominance in the global cable network. Figure 11 shows how different countries have contributed to the global expansion of cable and satellite networks over time. In the early years of submarine cables, the United Kingdom led the industry. It benefited from a first-mover advantage and strong access to credit, allowing British firms to build most of the early telegraph network. After WW2, however, British firms could no longer finance new systems at the same pace as they retired old ones. This led to a gradual decline in the UK's global role. During the same period, the United States expanded its network and emerged as the new dominant player.

Figure 12 focuses on the ownership shares of major powers. For cables, the turning point came with the fiber-optic revolution of the 1980s. American firms quickly adopted the new technology and rapidly expanded their networks, securing a dominant position in the global infrastructure. As with cables, the evolution of satellite ownership has been shaped by both geopolitical shifts and technological change. During the Cold War, the Soviet Union was the main player in the satellite segment. This was partly due to its geography: with limited sea access and a location closer to the poles, the Soviet Union had a comparative advantage in satellite-based systems. But after the Cold War, the collapse of the Soviet Union and Russia's economic crisis in the 1990s allowed the United States to compound its leadership in the cables segment with dominance in the satellites segments. Additionally, American firms strengthened their position again in the 2010s, by developing and rolling out new satellite systems based on technologies such as spot beam architecture and phased array antennas. These advances allowed the US to increase its market share steadily, even as new competitors such as Chinese companies entered the market and Russian as well

Figure 11. Ownership of telecommunications technology

(a) Ownership of the global submarine cable network

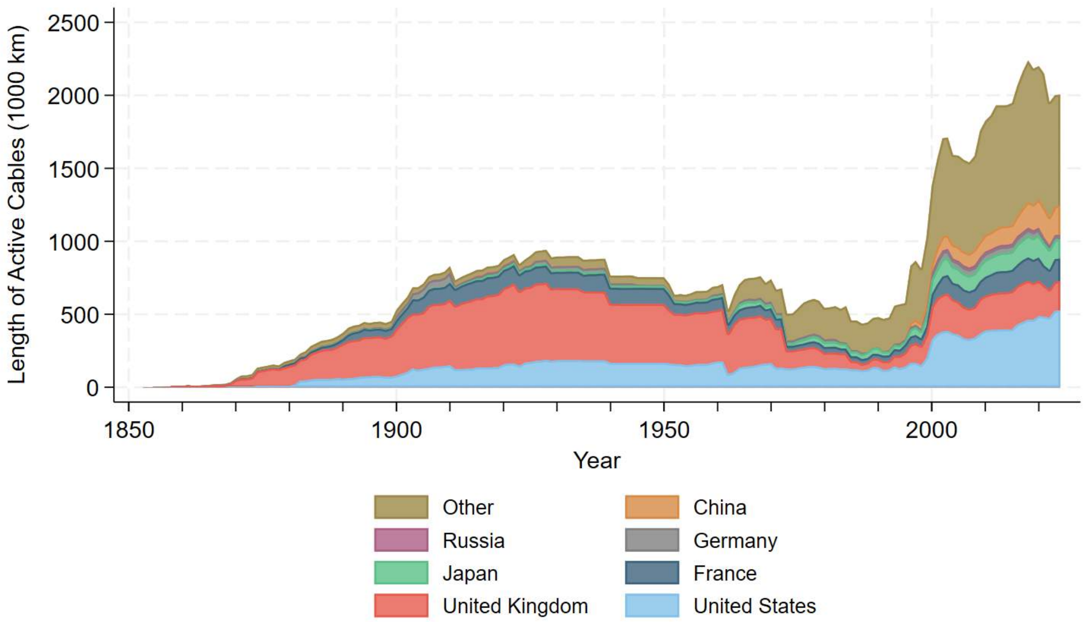

(b) Ownership of the global satellite network

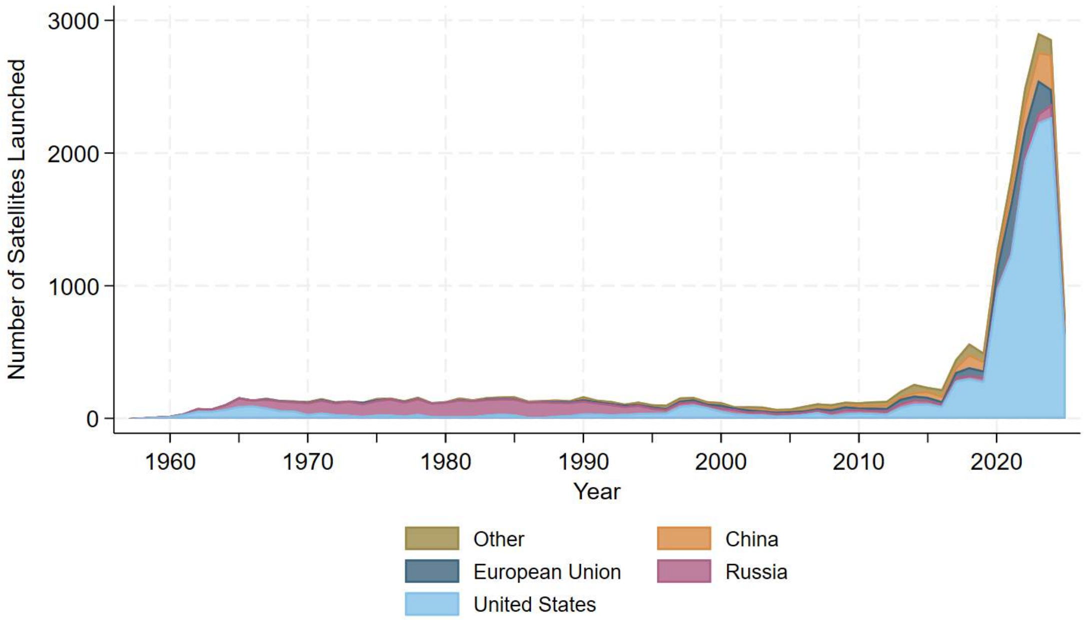

as European investments began to grow again.

Figure 12. Great powers' ownership of telecommunications technology

(a) Share of the global submarine cable network among great powers

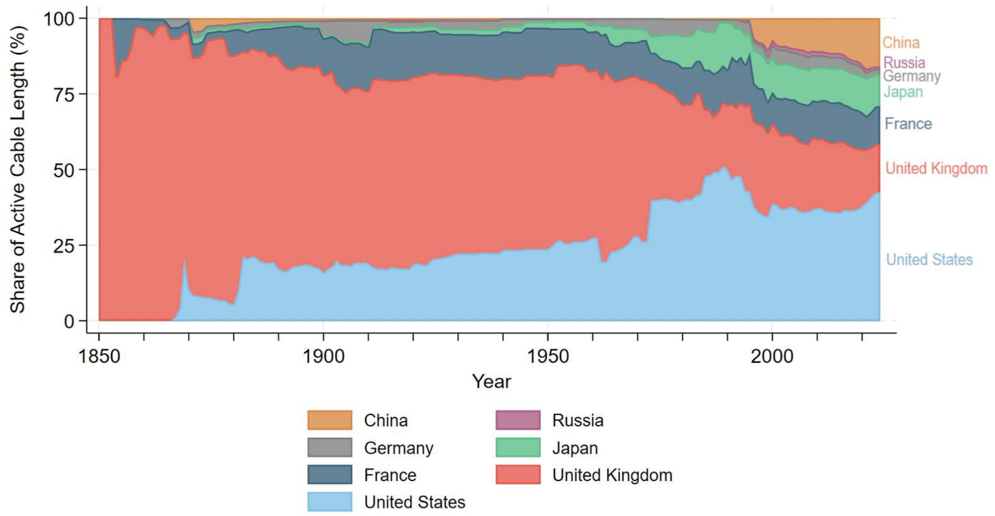

(b) Share of the global satellite network among great powers

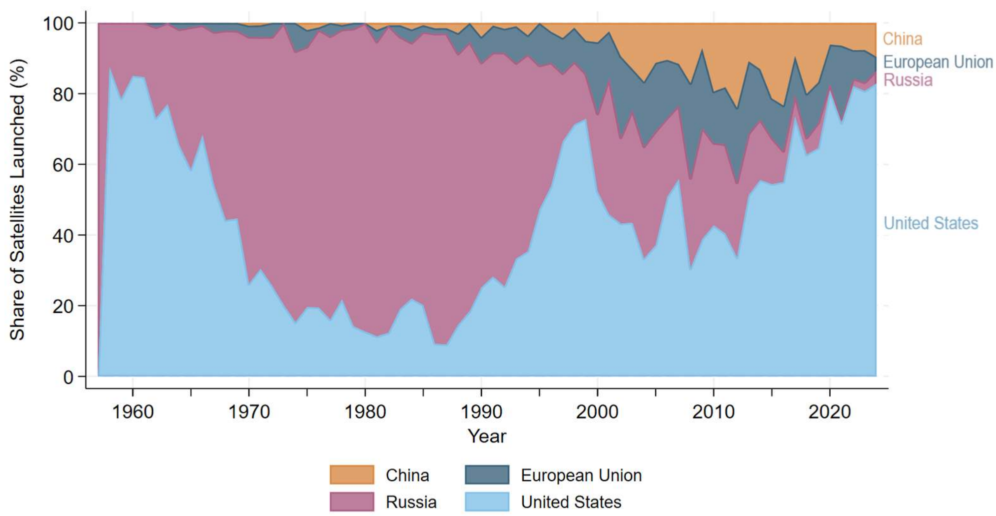

Hegemons as global backbone providers. Figure 13 compares the British-controlled global network in 1930 with the current US-controlled network. Cables that are fully or partly owned by UK or US firms are colored in red and blue, respectively. Both maps reveal a defining feature of hegemonic strategy: control over a genuinely global network that extends well beyond the hegemon's own borders or its regional sphere of interest. These networks serve not only to connect the hegemon's colonies, allies, and commercial partners, but also to provide backbone services to the rest of the world. Their wide geographical reach translates into both espionage capability and political leverage. It also creates strategic dependencies that extend well beyond ownership itself, as we discuss next.

Figure 13. Global reach of the hegemon's network: 1930 vs 2023

(a) 1930 Network: UK-controlled cables (in red, figure to be updated)

(b) 2023 Network: US-controlled cables (in blue)

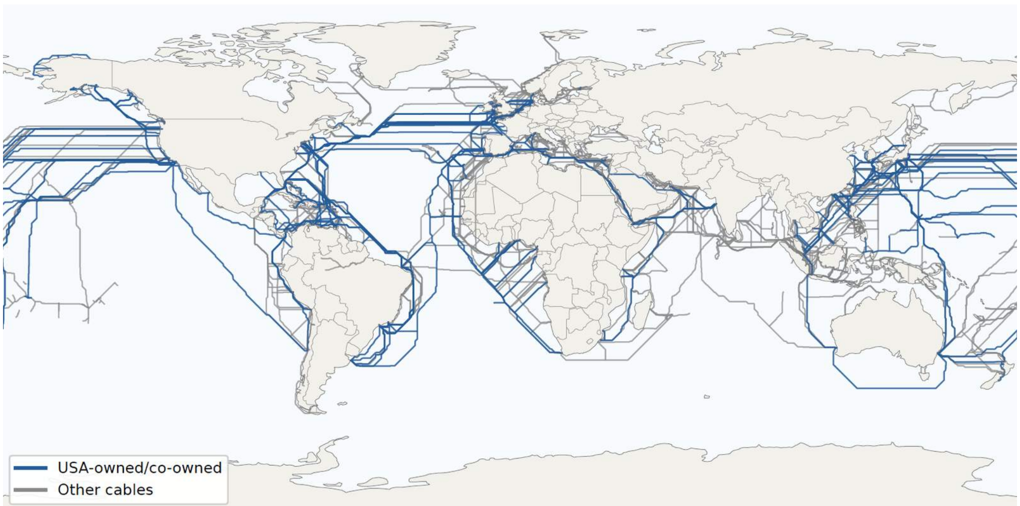

Hegemonic dominance is larger than commonly known. Figure 14 plots global dependence on US and Chinese cables using our model-derived dependence measure. Recall that dependence on a given country is the share of global data flows routed on cables owned, at least in part, by entities headquartered in that country. Dependence is therefore driven not only by the number of cables a country controls, but also by their geographic distribution and centrality in the network.

Figure 14. Share of global traffic routed on US vs. Chinese cables: 1990-2023

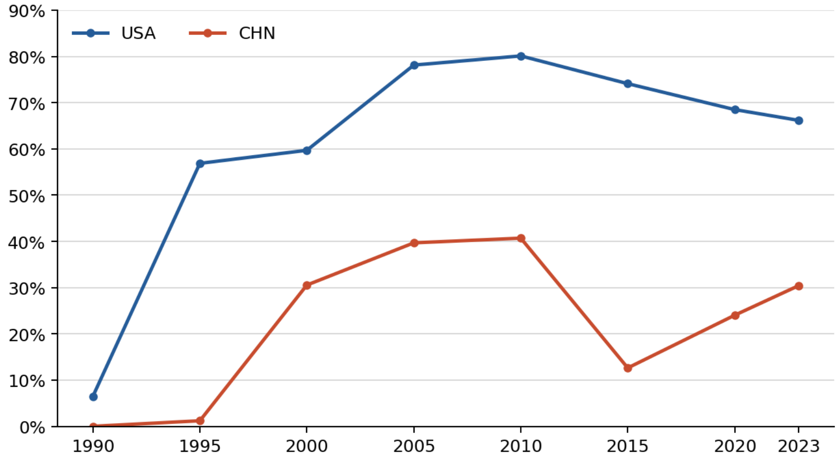

The dominance of the US is remarkable: around 70% of data flows traverse its cable network today according to our estimates, up from less than 20% in the early 1990s. This is further evidence that the roll-out of fiber-optic cables by US firms sharply increased US centrality in the global network. China is far less dominant globally, with 20 to 30% of data flows routed through its cables at some point. Historically, the UK comes close to matching US dominance today. As of 1900, we estimate that 65% of telegraph messages were routed through British-controlled cables. This confirms that hegemonic powers have been, and remain, the dominant players in digital infrastructure. Future versions of the paper will estimate historical dependency measures across time and space.

The global dependence numbers mask substantial geographic heterogeneity, however. A country is dependent on US or Chinese infrastructure when a large share of its traffic is routed on cables owned or co-owned by entities from that country.

Figure 15 shows that dependence on US-controlled infrastructure remains strikingly broad and global in 2023. It is most visible in the Western Hemisphere, but also high across several countries in Africa, East Asia, and Oceania, as well as parts of Europe and the Middle East. This reinforces our earlier finding that the United States acts as a global backbone provider with considerable leverage in the network.

The China map reveals a different but still consequential pattern. Dependence on Chinese-controlled infrastructure is more regionally concentrated, with especially strong exposure across East and Southeast Asia as well as the Gulf, and the Red Sea corridor. Relative to the United States, Chinese cable influence is less uniformly global, but in core theaters it is substantial.

Figure 15. Dependence on US and Chinese infrastructure in 2023

(a) Dependence on US cables, 2023

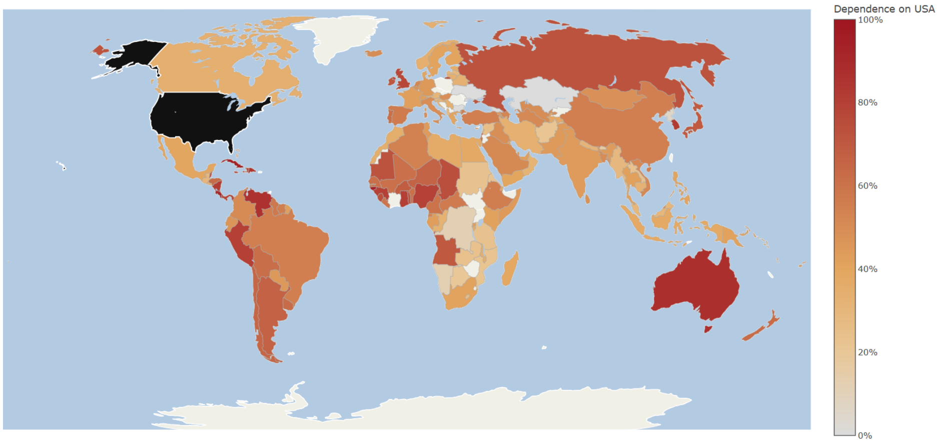

(b) Dependence on Chinese cables, 2023

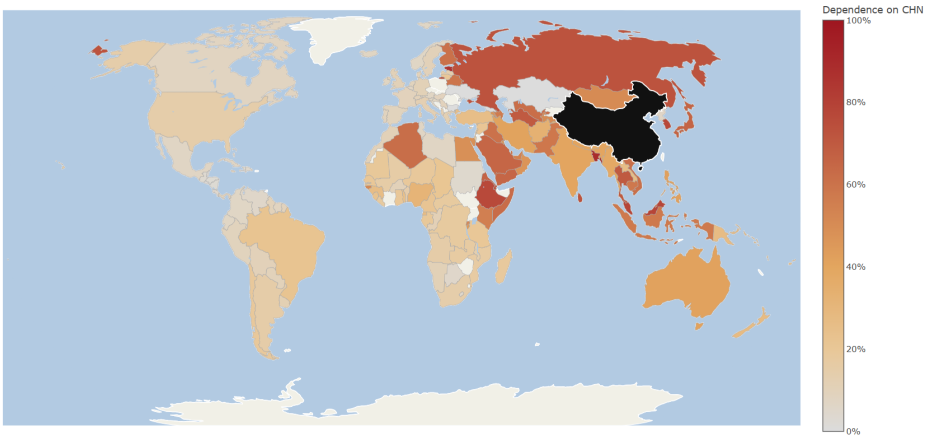

Note: These maps visualize country-level dependence on subsea cable infrastructure of a) the US and b) China. Darker shades indicate that a larger share of a country's traffic is routed on cables owned or co-owned by entities from the focal country. The focal country itself is shown in black.

So far, we measured dependence as the share of a country's traffic that is routed *through* a hegemon's cables. This is primarily a measure of exposure to surveillance. A country that carries another's data can intercept it, so the broad global reach of US cables documented above translates into a correspondingly broad capacity for interception and espionage. It does not, however, capture the welfare consequences of *denying* access to those cables.

As an alternative measure of dependency, we therefore turn to coercion and simulate an embargo in which a dominant owner bars a target country from using any cable it owns, forcing that country's traffic onto the remaining network. For this purpose, we use the full model toolkit and data from 2023.

Figures 16 and 17 report the resulting welfare losses for each potentially targeted country, for a US and a China cable embargo respectively. Two findings stand out. First, the estimated losses are substantial, which indicates that ownership dominance is not only an espionage advantage but also a source of coercive leverage. Countries that dominate digital infrastructure can impose measurable economic costs on others by withdrawing access.

Second, the welfare impact of an embargo is concentrated regionally. For the United States, losses are largest in Central and South America, where substitutes for US-owned infrastructure are scarcest. For China, the welfare losses from a cable embargo are lower overall, with the most affected countries being neighbouring Asian economies and some countries in Africa. In both cases, the welfare cost of an embargo is largest for the economies with the fewest alternatives.

Figure 16. Welfare loss from a US cable embargo

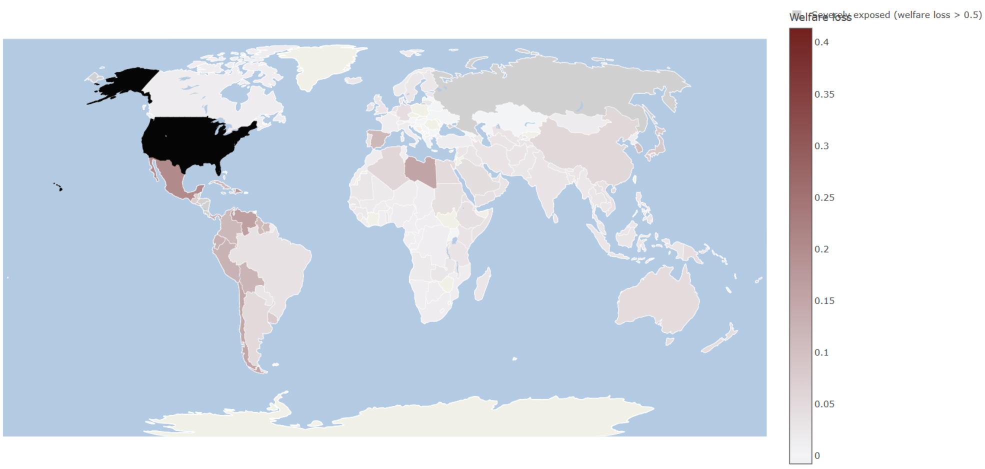

Note: The map shows the welfare loss per country in a scenario where that country is prohibited from using any cable with United States ownership, while all other countries retain full access. 2023 data.

Figure 17. Welfare loss from a Chinese cable embargo

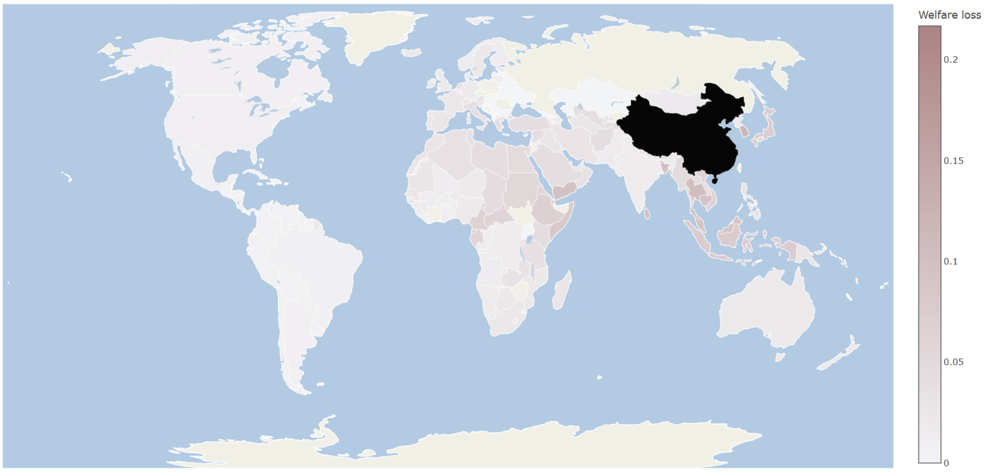

Note: The map shows the welfare loss per country in a scenario where that country is prohibited from using any cable with Chinese ownership, while all other countries retain full access. 2023 data.

6.2 Vulnerability: global data flows concentrate in a small set of (maritime) chokepoints

This section moves from dependence to the vulnerability of the subsea cable network. We study how exposed data flows are to systematic cable disruption, whether through classic methods such as anchor dragging or newer technology such as underwater drones and unmanned submarines. Vulnerability rises with the concentration of flows on a small number of cables and corridors, and falls with network redundancy and the scope for rerouting. Given the richness of the modern data on flows, capacity, and latency, we focus on the 2023 network, the most recent year for which all empirical components are available to estimate and validate the model.

Figure 18 maps estimated global data flows onto the full subsea cable network. At the cable level, flows are dispersed. A few cables carry high volumes according to our estimates, in particular those routed across the Atlantic and Pacific and along the Eurasian corridor through the Red Sea, but this partly reflects our shortest-path assignment in this step, which does not yet allow for rerouting.

Figure 18. Flows by cable: estimated global data traffic, 2023

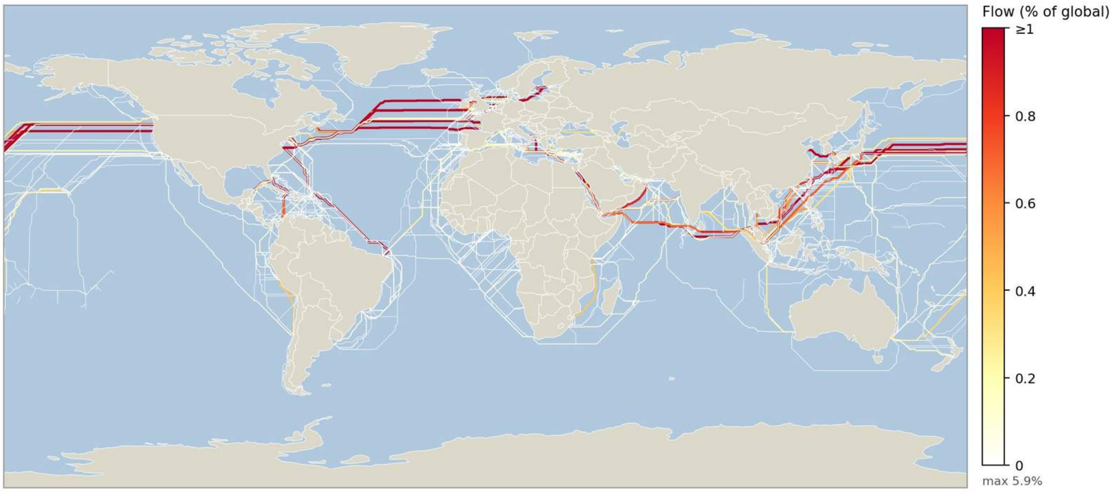

Notes: Each cable is colored uniformly along its entire route by the total traffic it carries, in percent of global data traffic. The legend is capped at 1 percent; the highest-flow cable carries 6.0 percent.

We next move from the cable level to geographic grid cells. This allows to better capture the spatial distribution of flows and identify particularly exposed areas. To do so, we overlay routed traffic onto a global two-degree grid, allocating each cable's traffic evenly across the cells it crosses. The grid cell allocation preserves total flows and sums to one hundred percent on the world map. Cell values can therefore be interpreted as the share of global data traffic that pass though this geographic area. This differs from the cable map, which paints each cable along its full length by its total flow. The grid instead concentrates traffic where many cables share the same narrow passage.

We then feed these cell-level exposures into the model to compute the welfare loss from disruption, accounting for rerouting and latency.

The grid approach is crucial for our purposes, as it reveals vulnerabilities and "digital chokepoints" that a plain cable map or landing-station data cannot. It shows how much traffic each area carries and which areas are most exposed to cable disruption in welfare terms. A grid approach can also inform infrastructure planning, in particular on where added redundancy would be most effective.

Figure 19 delivers a new, central result: trade chokepoints and digital chokepoints largely coincide. The same handful of passages that have constrained the movement of goods and naval power for centuries also constrain the movement of data. Traffic in the global telecom network concentrates on a small number of corridors, and these are, to a large extent, the world's maritime chokepoints,

Figure 19. Flows by grid cell: share of global data traffic, 2023

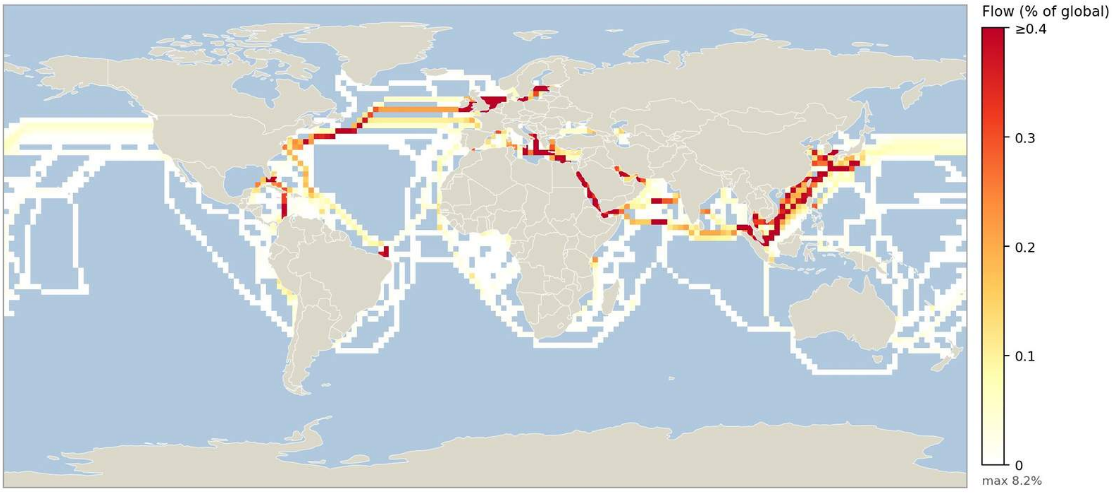

Notes: Each cable's traffic is split equally across the two-degree cells it crosses, so cell values are shares of global data traffic and sum to 100 percent worldwide. The legend is capped at 0.4 percent (the 95th percentile of cell shares); the most exposed cell, in the Irish Sea, carries 8.2 percent.

namely the Strait of Malacca, Suez, Hormuz, Gibraltar, and the English Channel. This is intuitive in hindsight, since cables and ships pass through the same narrow geography, but to our knowledge it has not been documented before.

The grid map identifies additional areas of concentration compared to the cable-level map in Figure 18. Collapsing flows by cell brings out the large-scale flows routed through the Mediterranean and Red Sea corridor, the Strait of Malacca, and the seas around China. It also reveals high flow concentration in regional chokepoints such as the Baltic Sea and the Strait of Hormuz, which are less visible in cable level maps. The open Atlantic and Pacific, by contrast, show less concentration than the cable-level assignment suggests, because transoceanic routes fan out over many parallel paths.

The magnitudes are striking. The most exposed grid cell, in the Irish Sea, alone accounts for 8.2 percent of global data traffic. Table 1 quantifies the exposure of the main corridors: for each, we count the cables crossing the corridor and the share of global traffic they carry. The English Channel and the British approaches are crossed by 44 cables carrying 25 percent of global data traffic; severing only the fifteen largest of them would cut links carrying 24 percent. Another 20 percent runs through the cable landing corridor on the US East Coast shelf between New York and Virginia, on just 32 cables. Even single grid cells channel remarkable volumes: one cell off New York is crossed by 18 cables carrying 10 percent of global traffic, and the cells at the Singapore

Strait and the Dover Strait each funnel close to 8 percent through an area of roughly 200 by 200 kilometers. The rapid advance of unmanned underwater vehicles makes coordinated cable cuts within such confined areas an increasingly plausible scenario.

Table 1. Digital chokepoints: cable exposure of maritime corridors, 2023

	Cables	Share of global traffic (%)	
Corridor	crossing	15 largest cables	All cables
English Channel / British approaches	44	24.4	25.0
US East Coast shelf (NY–Virginia)	32	19.8	20.4
Strait of Malacca / Singapore	29	9.2	9.7
Sicily Strait / central Mediterranean	32	7.9	8.1
Suez (Mediterranean approach)	13	7.7	7.7
Around Taiwan	17	6.7	6.7
Red Sea incl. Bab-el-Mandeb	17	6.5	6.5
Single grid cell off New York	18	–	10.4
Single grid cell, Singapore Strait	21	–	7.6
Single grid cell, Dover Strait	8	–	7.6

Notes: Each corridor is a compact geographic window around the named passage (single two-degree cells in the bottom panel). Cables are counted once each, at the flow they carry through the corridor; the last two columns report the combined traffic of the fifteen highest-capacity crossing cables and of all crossing cables, as a share of global data traffic in 2023. Cutting a corridor severs these links and forces global rerouting; the welfare consequences are shown in Figures 20 and 21.

We now turn from traffic to welfare. Figure 20 shows the global welfare loss from removing each cable individually, accounting for rerouting. The picture is far more benign than the raw flow map suggests.

The welfare impact of disrupting a single cable is close to zero for most cables, and the map is nearly white. Some cables show up as more exposed than others, in particular along the Eurasian corridor through the Red Sea and the Indian Ocean, but overall, for individual cables, disruption risks and global welfare losses are small. This is not surprising. The internet is built for resilience, and operators successfully manage individual cable failures on a daily basis. From a geoeconomic perspective, the more interesting scenario is coordinated cable cutting or wartime damage across a broader area.

We therefore repeat the welfare computation at the grid-cell level. You can think of this exercise as a thought experiment in which every cable crossing that cell is severed at once, and we re-solve the model to obtain the global welfare loss, while accounting for rerouting. Intuitively, the resulting welfare number by grid

Figure 20. Global welfare loss by cable, 2023

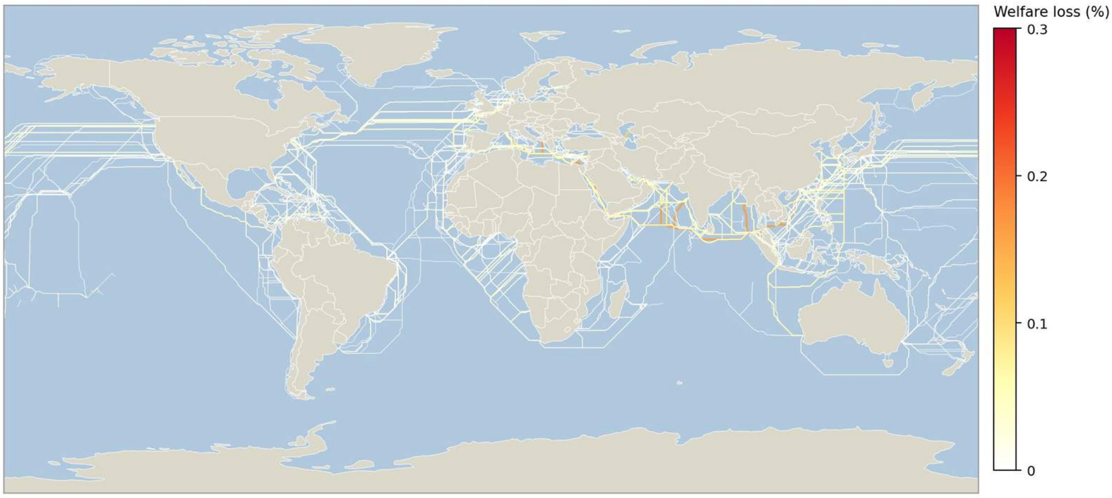

Notes: Welfare loss from removing each cable individually, accounting for rerouting, in percent; 2023 network. The color scale is identical to the grid-cell welfare map in Figure 21, so the two maps are directly comparable.

cell captures what global welfare losses if this square of ocean went ”dark”. Such an event could in principle arise from natural causes such as a major earthquake, but the setting we have in mind is military tension or open conflict, in particular a coordinated sabotage campaign or collateral damage when ships sink or bombs fall. The cell is thus treated as the unit an adversary would attack, rather than the cable. Large-scale cable disruptions of this sort would likely require long repairs given the limited global supply of repair ships. In an active war zone, repairs may be very difficult or impossible altogether, extending the outage well beyond the one to three weeks it usually takes to fix a single cable.

The welfare results in Figure 21 point to a large impact of coordinated cable ruptures in a broader area. Vulnerability is highest in classic maritime chokepoints such as the Strait of Malacca and the Red Sea–Suez corridor, but also on the North Atlantic plateau east of Newfoundland. As expected, the welfare map looks markedly different from the map of raw data flows in Figure 19, as we now account for rerouting, congestion, latency, and economic linkages. The usual maritime chokepoints show up prominently in both maps, but the welfare-based estimates show little vulnerability in the East and South China Seas and the Pacific. In these areas, a large number of geographically dispersed cables allows for quick rerouting, implying less vulnerability to cable cutting at the grid level.

Figure 21. Global welfare loss by grid cell: vulnerability to disruption, 2023

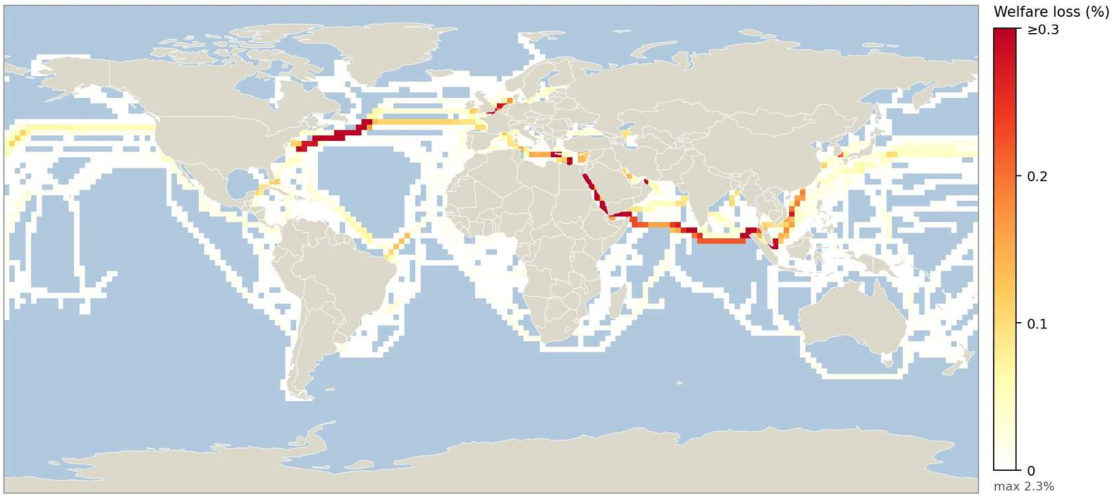

Notes: Welfare loss from severing all cables crossing each cell at once, accounting for rerouting and latency, in percent. The color scale is identical to the cable-level welfare map in Figure 20.

6.3 Cable wars: historical evidence and modern simulations for Taiwan and the South China Sea

This section extends the preceding vulnerability analysis to worst-case scenarios in which we assume that entire countries are disconnected by coordinated cable cutting, or in which whole segments of the global network are taken out at once.

6.3.1 Historical evidence on coordinated cable cutting: 1898 and 1914

History shows that such scenarios are far from unprecedented. The best-known case of a complete cable-cutting embargo is Cuba during the Spanish-American War. As discussed in Section 2, the United States cut or censored every cable link between Cuba and Spain, and this proved a decisive advantage in military communications and in the conduct of the war. A US Navy officer later concluded that cables "exercise a dominating influence upon the control of the seas" and that it was mostly "coal and cables" that won the war (Winkler, 2008). Figure 22 illustrates this blockade, showing all telegraph cables around Cuba as of early 1898. The red crosses mark the affected cables rather than the precise points of severance. The cable to Florida was censored at the US landing point in Key West.

A second case is WW1, when Britain severed most of Germany's direct overseas cables and later cut its rerouted connections in the Atlantic. Figure 23

Figure 22. US cable cutting in the Spanish-American War of 1898

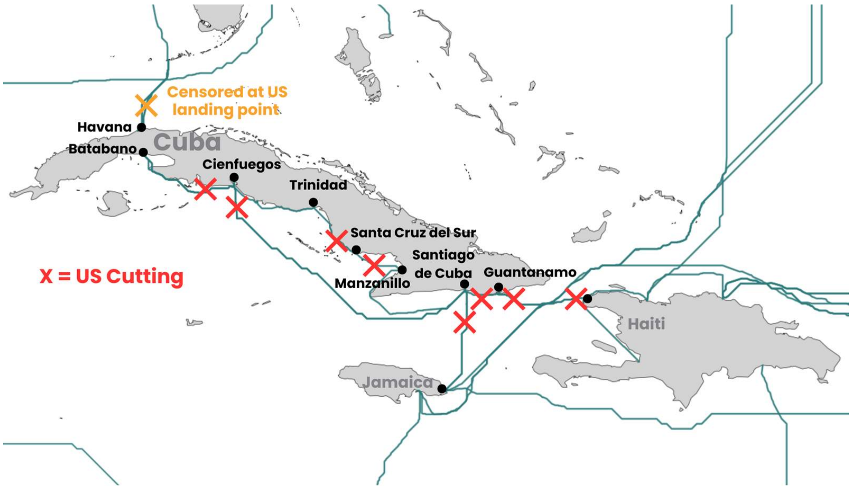

Figure 23. Cable cutting during WW1 (1914/1915)

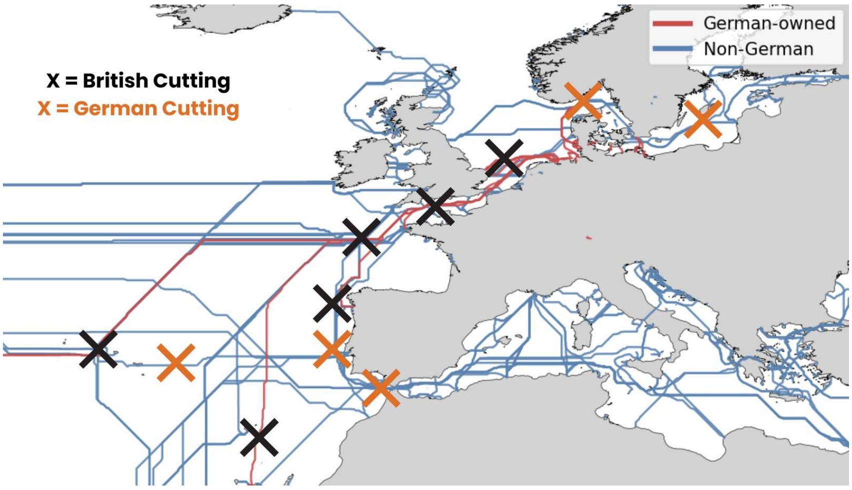

shows the telegraph network around Europe as of early 1914. Red cables are German-controlled (including state-owned) and blue cables are controlled by other powers. The black and brown crosses mark the targeted cables rather than the precise points of severance. The figure is stylized and does not include every cut. Appendix D provides a systematic account and shows that our network-based vulnerability measure predicts which cables were targeted.

These cuts isolated Germany from global communications and enabled British espionage, since any German telegram now had to travel through enemy cables. The most consequential result was the interception of the Zimmermann Telegram, in which Germany proposed a military alliance with Mexico against the United States. Its decryption and publication in early 1917 caused a public outcry in the US and reportedly accelerated the American decision to enter the war.

6.3.2 Simulating coordinated cable cutting today: Taiwan and the South China Sea

What would a similar scenario look like today? Using our model, we simulate two disruption cases. First, the isolation of Taiwan by cutting all of its international cables. Second, a WW1-style cable war in the South China Sea, in which both Chinese and US-controlled cables are severed en masse.

Isolating Taiwan: Taiwan connects to the world through 14 international undersea cables, which carry more than 99 percent of its external data traffic (TWNIC, 2025). Satellites play only a marginal role due to limited capacity. Government reports put satellite capacity at roughly 2 Gbps in total, compared to the cable system's total transmission capacity of about 83 Tbps, i.e. only about 0.0025 percent of cable capacity (MODA, 2023). Over the past years, the Ministry of Digital Affairs (MODA) rolled out more than 700 non-geostationary satellite ground stations across Taiwan and its outlying islands, drawing on OneWeb (LEO) and SES (MEO) capacity. This network is designed for emergency government command traffic, mainly messaging and low-rate video, rather than civilian or economic demand (MODA, 2024).

The alternative Starlink system, the largest low earth orbit satellite constellation, does not serve Taiwan as of mid-2026. Negotiations broke down over Taiwan's legal requirement that a foreign operator form a joint venture with a local partner holding majority ownership, which SpaceX refused (Hsu, 2026; Taipei Times, 2026). SpaceX is also viewed as compromised by Elon Musk's business interests in China, and Musk has acknowledged that Chinese officials told him not to launch Starlink in China, a position Beijing extends to Taiwan (Hsu, 2026). In addition, SpaceX treats Taiwan as a low priority market, citing 4G and 5G coverage above 99 percent (Taipei Times, 2026).

Even full access to Starlink would not change the capacity picture, however. Starlink delivers bandwidth per satellite, and only the satellites passing over a given area serve it. Given Taiwan's small size, just a handful of satellites cover the island at any moment, so the instantaneous Starlink capacity over Taiwan is on the order of a few hundred Gbps at most under current hardware. This remains well below 1 percent of the roughly 83 Tbps carried by Taiwan's cables. Satellites, in other words, cannot substitute for more than a negligible share of cable capacity.

To be conservative, we nonetheless assume in the simulation that Taiwan has access to Starlink, which it currently does not, and that satellites carry 1 percent of its pre-sabotage cable traffic, a generous upper bound.

Figure 24 shows the results of the simulation, in which all 14 of Taiwan's international cables are cut within a short window. More specifically, we assume that all subsea cable segments in the Taiwan Strait, or within 100 km of Taiwan, are disabled, a total of 20 distinct systems. Disabled cables in this scenario are marked in black, while country shading gives the resulting welfare loss.

For Taiwan's neighbors the loss is limited, as a result of existing redundancy. Data can be rerouted onto the many unaffected cables in the region at comparative low cost. For Taiwan itself, however, the loss is severe. Even under the generous assumption on satellite access, the welfare shock is large, because rerouting is no longer available and the island goes almost entirely dark in terms of bandwidth and real-time connectivity.

Figure 24. Isolating Taiwan - welfare effects and affected cables

Note: Disrupted cable systems are shown in black and country shading gives the welfare loss by country. For comparability with Figure 25 and to better show the variation across impacted countries with low welfare losses, the color scale is capped at 0.6, with values above that uniformly colored as 0.6. Taiwan's own welfare loss exceeds 2.2.

A cable war in the South China Sea: Figure 25 simulates a larger regional cable war, in which cables passing the South China Sea are severed by the warring parties. Specifically, we assume that 82 cable systems with 209 cable segments are damaged. This scenario echoes the cable-war theaters of WW1 and WW2, where dozens of cables were cut. The global welfare impact is far broader than in the Taiwan case. Losses are large across the neighboring region and spill over well beyond it, reaching countries in Africa and the Near East that rely on this corridor for internet connectivity. Even under generous assumptions of Starlink satellite backup, the welfare cost of a cable war in the South China Sea is severe.

Figure 25. Cable war in the South China Sea - welfare effects and affected cables

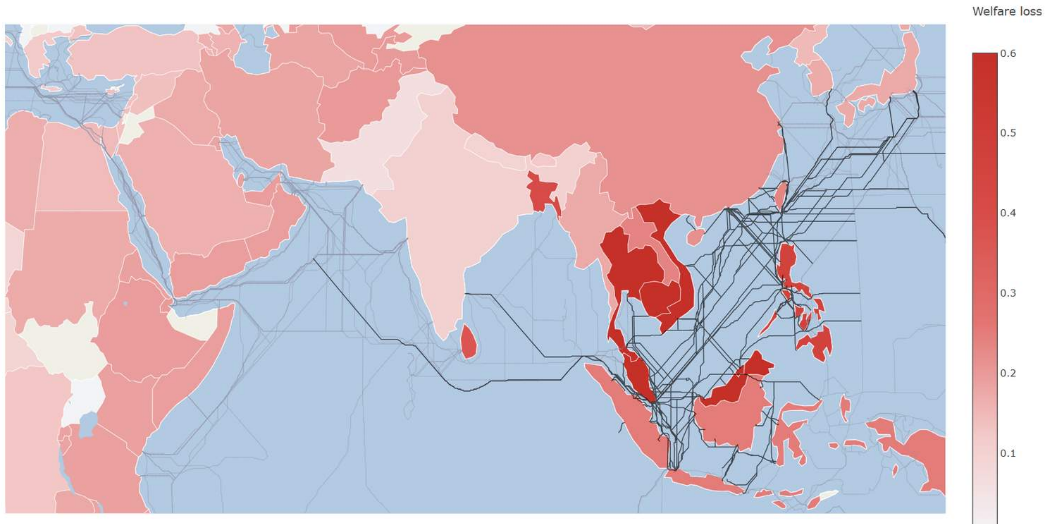

Note: This map shows the welfare loss incurred per country in a scenario where all cable segments in the South China Sea are disabled. For comparability with Figure 24 and to better show the variation across impacted countries with low welfare losses, the color scale is capped at 0.6, with values above that uniformly colored as 0.6.

6.4 Geopolitical risk and cable investment: decoupling and chokepoint avoidance

How do telecom firms and governments respond to growing geopolitical risks and the weaponization of telecom networks? This section leverages our long-run, granular *GlobeTel* dataset to study investments in new cables in history and today. The main focus is on strategic decoupling and chokepoint avoidance, as we want to understand whether investment decisions are shaped by geopolitical tensions and weaponization events.

We start with a case study of the "weaponization" shocks of 1898/99, a two-year period in which telegraphy cables became prominent targets in conflict (especially during the Spanish American War of 1898) and in which Britain started to openly weaponize its dominance in cable networks at the expense of rivals (Fashoda incident of 1898, and Britains's 1899 cable embargo during the Second Boer War). Both the cable-level data and archival sources such as parliamentary speeches strongly suggest that these weaponization shocks played a decisive role in motivating the cable building frenzy of the early 20th century. We document that, between 1899 and 1913, France, Germany and the US considerably expanded their own cable systems, often on the same corridors dominated by Britain. This evidence is corroborated with a set of regressions, which aim to disentangle economic motives of cable building (market size and entry costs) from security motives more rigorously. In a supplementary step, we draw on newly collected data on new and *planned* cables, covering all recently finished cable systems and those aimed for completion up until 2035. The routing characteristics of these new cable projects show that chokepoint avoidance has become an important design principle of the ongoing investment wave, a fact supported by industry reports that emphasize the role of geopolitics as driving cable investments in general, and routing choices, in particular.

The key take away is that militarized tensions, above all with the hegemon, are associated with decoupling investments, in which countries build their own redundant capacity to decouple from networks they do not control. The same logic is visible today as new and planned cables are routed to avoid geographies dominated by rival powers, eg. the South China Sea, as well as other contested chokepoints.

6.4.1 Decoupling from the hegemon: evidence from history

Case study: cable investments after the 1898/99 shocks. The years 1898 and 1899 were decisive for global telegraphy, rapidly politicizing a network previously seen as a largely commercial arena. The seminal book by Headrick (1991, p.7) summarizes this episode as follows: "Cable cutting in the Spanish-American War and censorship in the Boer War showed the world how vulnerable were the lines of communication, and how much more vulnerable were the nations whose global communications depended on the goodwill of Great Britain. In reaction, France and Germany hastened to build their own networks, with heavy government subsidies [. . .]." Here, we bring Headrick's "decoupling" hypothesis to the data for the first time, by studying the dynamics of cable investments before and after the 1898/99 shocks. Figure 26 plots the kilometers of new cable laid each year by non-UK powers between 1885 and 1913 by country, using our dataset. Before 1899 non-British cable building is low and irregular. After 1898, it rises rapidly and stays elevated for a decade. Germany, which had laid almost no submarine cable before 1899, rapidly expands its own cable network, while the US and France expand substantially as well.

Figure 27 shows where the new capacity went, splitting the same non-UK construction by the share of UK ownership pre-1898. It shows that a rising share of post-1898 construction by Britain's rivals flows into corridors it had previously dominated. Simultaneously, construction of new routes picks up, which is mainly driven by the United States investing into new Pacific routes to move away from the contested Atlantic ones.

Figure 26. Non-UK cable building, by country, 1885 - 1913

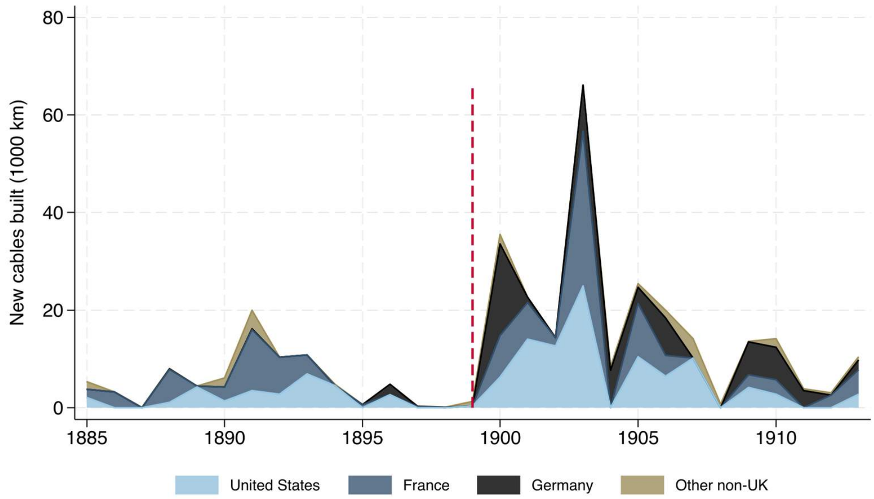

Notes: The figure shows the kilometers of new submarine cables laid each year by non-UK owners from 1885 to 1913, stacked by owner country (United States, France, Germany, and all others). Cables are counted in the year they enter service. The dashed line marks 1899.

Archival evidence and a large body of historical research supports the view that the observed patterns in the data were deliberate and centrally coordinated by government, which aimed to decouple from Britain and build an own, independent cable network. In Germany, the Kaiser and its government launched a new cable policy in 1899/1900. In a letter to Chancellor von Bülow from October 1900, postal minister Victor von Podbielski describes plans for a "worldwide German cable network" that would free the Reich from dependence on Britain (Wroblewski, 2016, p. 10). The central mechanism towards that goal were substantial government subsidies and state funding. "In November 1899 the government asked the Reichstag for credits to build a cable network" (Headrick, 1991a, p. 106). It also meant the built-up of an own cable industry. The state-backed Deutsch-Atlantische Telegraphengesellschaft was founded in the same year and opened a direct Emden to New York cable in 1900. Several other state-linked companies followed.

France pursued a similar strategy, and heavily relied on government funding

Figure 27. The new cables are largely built in corridors dominated by the UK

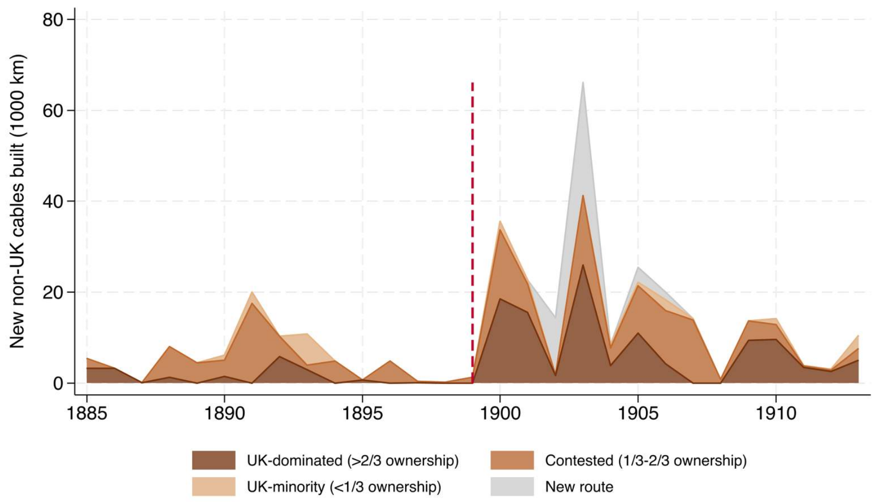

Notes: The figure shows non-UK new submarine cable laid each year from 1885 to 1913, split by how dominant the UK was in a corridor before 1898. A corridor's UK share is the UK-owned fraction of all cable laid in that corridor as of 1897. Corridors are labelled "UK-dominated" when this share exceeds two thirds, "contested" between one and two thirds, and "UK-minority" below one third. "New route" denotes corridors with no cable before 1898. Cables are counted in the year they enter service. The dashed line marks 1899.

and subsidies. A 1899 resolution before the Chamber of Deputies, backed by the ministers of war, navy and colonies, declared it "necessary for a great country like France to have independent national telegraphic communications" (Headrick, 1991a, p. 103). The United States concentrated new investments on the Pacific, after acquiring the Philippines, Hawaii and Guam in 1898. President McKinley urged Congress that the new cable system to these new territories be "wholly under the control of the United States, whether in time of peace or of war" (McKinley, 1899), a commitment realized in the Commercial Pacific Cable Company a few years later. Britain, for its part, also continued to invest heavily and expand its network, completing the state-backed "All-Red Line" in 1902, a system of cables around the world which interconnected the Empire without touching foreign soil.

Decoupling from the hegemon: regression evidence. We next bring the decoupling hypothesis to the data more systematically, moving beyond case study evidence. To do so, we use cable investments as dependent variable in a Poisson Pseudo-Maximum Likelihood (PPML) framework and draw on the full *Globe-Tel* dataset, 1850-2025. We think of the construction of a new submarine cable as entry into a corridor-specific transit market. Standard economic reasoning suggests that entry into a corridor should increase with the size of the market it serves (e.g., expected data flows) and fall with cost of entry. However, if firms and states also value reducing their dependence on adversarial powers, cable construction should in addition respond to the geopolitical climate and rise when international tension is high. To distinguish economic and geoeconomic motives, we start by estimating the following corridor (m)-year (t) PPML regression:

$$\begin{aligned}\text{NewCablesBuilt}_{mt} = {} & \beta\, \text{GlobalTraffic}_{mt} + \gamma\, \text{Threats}_t \\ & + \alpha_m + \delta_m\, t + \text{Controls}_{mt} + \varepsilon_{mt}\end{aligned}$$

The dependent variable $\textit{NewCablesBuilt}_{mt}$ is the total kilometers of new cables brought into service in corridor m and year t. $\textit{GlobalTraffic}_{mt}$ is the log of expected traffic routed through the corridor, our measure of market size. $\textit{Threats}_t$ is a five-year moving average of the number of ongoing militarized interstate disputes from Correlates of War Militarized Interstate Dispute (MID) data (version 5), which records explicit threats and displays of force between states back to 1816. The moving average smoothes noisy annual counts and lets construction respond to tension with a lag. We include corridor fixed effects α_m (absorbing time-invariant differences in entry cost) and corridor-specific linear trends $\delta_m t$. $\textit{Controls}_{mt}$ contains the lagged corridor cable stock (in logs), which proxies market saturation, and global trade openness (*GlobalTrade*/*GlobalGDP*). Because $\textit{Threats}_t$ varies over time but not across corridors, γ is identified from the co-movement of building across corridors with global threats. A positive γ indicates that cable building rises with geopolitical tension beyond economic incentives.

Table 2 reports our aggregate-level results. As expected, market size matters. The elasticity of building with respect to global traffic is about 0.3 to 0.5, so a 10% increase in corridor traffic is associated with 3 to 5% more new cable building, and construction is weaker where the existing cable stock is larger, consistent with entry costs being higher due to more competition. More interestingly, geopolitical threats have a positive and significant coefficient. One additional worldwide threat is associated with about 6% more new cable in column (1) (i.e. $e^{0.058} - 1$), or roughly 60% for a one-standard-deviation increase in tension. If we restrict threats to only those issued by the hegemon (the UK before 1939, the US from 1939), the association with cable building becomes stronger with about 13% per additional hegemon-issued threat in column (2) (roughly 20% for a one-standard-deviation move). Broadening the threat concept to include disputes that escalated to war leaves a positive, significant coefficient in column (3), so the results are robust to the definition of tension. The threat coefficients are also stable across different moving-average windows. The aggregate results show

Table 2. Effect of Geopolitical Threats on Corridor-Level Investment

	(1)	(2)	(3)
	$NewCablesBuilt_{mt}$		
Global Traffic (log)	0.318***	0.461***	0.299***
	[0.067]	[0.097]	[0.076]
Global Threats	0.058***		
	[0.009]		
Hegemon-Issued Threats		0.126**	
		[0.063]	
Global Threats + Wars			0.035***
			[0.006]
Lagged Cable Stock (log)	-0.089*	-0.069	-0.071
	[0.047]	[0.047]	[0.047]
Corridor FE	Yes	Yes	Yes
Observations	2,703	2,703	2,703

Notes: Dependent variable: kilometers of new cable built in corridor m and year t. *Global Traffic* is the log of our estimate of all traffic flows routed in corridor m in year t, based on trade flows and our traffic-assignment algorithm. *Global Threats* is the five-year moving average of the number of militarized disputes (at the level of a threat or display of force) ongoing worldwide in a given year, from the Correlates of War MID database. *Hegemon-Issued Threats* restricts this to disputes in which the hegemon (the UK before 1939, the US from 1939) issues the threat. *Global Threats + Wars* additionally counts disputes that escalate to the use of force or war. *Lagged Cable Stock (log)* is the log of one plus the total kilometers of cable in the corridor at year $t-1$. All columns include corridor fixed effects and corridor-specific linear time trends, and condition on global trade openness *(GlobalTrade/GlobalGDP)*. War years (1914 - 1919, 1939 - 1945) are excluded. Robust standard errors clustered at the corridor level in brackets. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

that cable building co-moves with global threats, but it treats every country symmetrically. We next disaggregate to the country (j)-corridor (m)-year (t) level, which allows us to identify differential effects of geopolitical threats across countries and to include bilateral measures of tensions. The dependent variable in this setting is kilometers of new cable built by country j in corridor m and year t, which we again estimate by PPML:

$$\begin{aligned}\text{NewCablesBuilt}_{jmt} = {} & \beta_1 \, \text{GlobalTraffic}_{mt} + \beta_2 \, \text{DomesticShare}_{jmt} \\ & + \gamma_1 \, \text{GlobalThreats}_t + \gamma_2 \, \text{BadRelations}_{jt} \\ & + \gamma_3 \, (\text{GlobalThreats}_t \times \text{BadRelations}_{jt}) \\ & + \alpha_{jm} + \delta_m t + \text{Controls}_{jmt} + \varepsilon_{jmt}\end{aligned}$$

GlobalThreats$_t$ is the measure of worldwide militarized tension from Table 2, and *BadRelations*$_{jt}$ is a dummy equal to one for countries above the median in foreign-policy distance to the hegemon, based on the Correlates of War foreign policy similarity index. The coefficient of interest is the interaction γ_3, which measures whether countries with bad relations with the hegemon construct more than others when global tension is high. *GlobalTraffic*$_{mt}$ again measures the size of the corridor market and *DomesticShare*$_{jmt}$, country j's own traffic divided by global traffic in the corridor, proxies the corridor's strategic value to j. Corridor-country fixed effects α_{jm} absorb any time-invariant propensity of country j to build in corridor m and $\delta_m t$ are corridor-specific trends. The remaining controls are the lagged corridor cable stock, which captures saturation from past investment, and country-year macro controls (log GDP, log population, GDP per capita). Because *BadRelations*$_{jt}$ varies by country and year, γ_3 is identified within corridor-country cells.

Columns 1 to 3 in Table 3 show that global threats are positively and significantly associated with increased cable building, especially for countries non-aligned with the hegemon. For countries aligned with the hegemon, a one-standard-deviation rise in worldwide threats is associated with about 47% more new cable building. The interaction term is positive and statistically significant, indicating that countries with bad relations with the hegemon, build roughly 75 percentage points more per one-standard-deviation rise in global threats, on top of the 47% baseline[10]. The results are similar when macro controls and war years are added (columns 2 and 3).

Columns 4 to 6 use a bilateral measure of geopolitical tensions: *Hegemon Threats* is the five-year moving average of the yearly number of outright threats and militarized disputes between country j and the hegemon, again using the MID database. We find that bilateral threats vis-a-vis the hegemon are positively and significantly associated with a country's own cable building. A

[10] $e^{0.046\times 8.3} - 1 \approx 47\%$ for countries with good relations with the hegemon versus $e^{(0.046+0.050)\times 8.3} - 1 \approx 122\%$ for those with bad relations.

one-standard-deviation increase in bilateral threats is associated with about 20% more construction, and the estimated relationship is stable as macro controls and war years are added.

Across both specifications the lagged cable stock enters negatively, consistent with corridor saturation reducing new investment, and the domestic traffic share enters positively.

Taken together, the regression analysis results are consistent with the decoupling hypothesis. Geopolitical tensions are associated with more cable building and the response to geopolitical threats appears most pronounced among those countries in rivalry with the hegemon.

6.4.2 Chokepoint avoidance and decoupling today: evidence from new and planned cables

In recent years, history appears to be repeating itself. A June 2026 article in The Economist argues that growing geopolitical risks have prompted firms and governments to launch new cable systems that avoid maritime chokepoints and increase resilience against conflict (The Economist, 2026). Our collected data on newly laid and planned cables are consistent with this view: chokepoint avoidance has become an important feature of the ongoing cable investment wave (see also Mauldin, 2024).

Prominent examples include the Echo and Bifrost systems, developed by Meta and Google with regional partners and in service since 2025. These cables were rerouted to circumvent the South China Sea through Indonesian waters, as was the Apricot system that is also backed by Google and Meta (Murphy and Bryja, 2025). These rerouted paths are longer and costlier than the direct routes they replaced, indicating a willingness to pay large premia for chokepoint avoidance (Murphy and Bryja, 2025). Three additional major cable systems are under construction or in advanced planning and are likewise designed to avoid chokepoints: Blue-Raman, owned by Google and Telecom Italia Sparkle, will connect Europe to India while bypassing Egypt, the currently dominant chokepoint corridor (Carnegie Endowment, 2024); Meta's Project Waterworth, a multibillion dollar system spanning some 50,000 km, will be routed to circumvent both the Red Sea and the South China Sea (IEEE, 2025); and Far North Fiber, a joint venture of Finnish, US, and Japanese network operators, will connect Europe to Japan through the Arctic Northwest Passage, avoiding the southern chokepoints altogether (Middleton and Rønning, 2022).

Current government financing initiatives follow the same logic. The Euro-

Table 3. Effect of Geopolitical Threats on Country-Corridor Investment

	(1)	(2)	(3)	(4)	(5)	(6)
			NewCablesBuilt$_{\text{jmt}}$			
Global Traffic (log)	0.176	0.057	0.053	0.325**	0.151	0.139
	[0.112]	[0.155]	[0.109]	[0.128]	[0.226]	[0.171]
Domestic Traffic Share	7.370***	4.502**	4.598***	6.972***	3.200	3.673**
	[2.014]	[2.044]	[1.770]	[1.368]	[1.974]	[1.560]
Global Threats	0.045***	0.046***	0.046***			
	[0.016]	[0.015]	[0.015]			
Bad Relations	-0.253	-0.110	-0.136			
	[0.419]	[0.513]	[0.511]			
Global Threats × Bad Relations	0.054***	0.050***	0.050***			
	[0.011]	[0.011]	[0.010]			
Hegemon Threats (bilateral)				0.743***	0.701***	0.716***
				[0.172]	[0.254]	[0.265]
Lagged Cable Stock (log)	-0.080*	-0.111**	-0.113**	-0.062*	-0.105**	-0.109**
	[0.044]	[0.054]	[0.052]	[0.035]	[0.051]	[0.049]
Corridor-country FE	Yes	Yes	Yes	Yes	Yes	Yes
Corridor-specific trends	Yes	Yes	Yes	Yes	Yes	Yes
Macro controls	No	Yes	Yes	No	Yes	Yes
War-year dummies	No	No	Yes	No	No	Yes
Observations	12,691	12,611	13,610	12,691	12,611	13,610

Notes: Dependent variable: kilometers of new cable built by country j in corridor m and year t. Global Traffic is the log of estimated traffic routed in corridor m and year t, a measure of the corridor's profitability. Domestic Traffic Share is country j's own traffic divided by total traffic in the corridor. Hegemon Threats is the five-year moving average of the number of militarized disputes (at the level of a threat or display of force) between country j and the hegemon (UK before 1939, US from 1939) that are ongoing in a given year. Global Threats is the worldwide bilateral threat count. Bad Relations is an indicator for above-median foreign-policy distance from the hegemon, using the Correlates of War foreign-policy similarity index. Lagged Cable Stock (log) is the log of one plus the total kilometers of cable in the corridor at year $t-1$. Macro controls are log rGDP, log population, and rGDP per capita. The sample is the twelve major cable-owning powers, the UK and the US are excluded when they are the hegemon, since the threat and policy distance measures are undefined for the hegemon itself. War years are excluded except in columns 3 and 6, which retain them and add war-year dummies for the WW1 (1914 - 1919) and WW2 (1939 - 1945) periods. All columns include corridor-country fixed effects and corridor-specific linear time trends. Robust standard errors clustered at the corridor and country level in brackets. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

pean Commission's 2026 list of Cable Projects of European Interest includes an additional Arctic connection to Asia via two new publicly supported routes, with an estimated public-private cost of EUR 4.3 billion (European Commission, 2026, p. 34). The EU's expert group explicitly justifies this investment by the fact that "90% of data traffic between Europe and Asia transits through submarine cables located in the Red Sea" (European Commission, 2026, p. 34), which is in line with independent industry estimates (Mauldin, 2024). In a similar vein, Russia's state owned Polar Express cable, under construction from Murmansk to Vladivostok, runs entirely within Russian jurisdiction with the aim to increase digital sovereignty and insulation from foreign interference (Middleton and Rønning, 2022).

Government interventions are also increasingly aimed at strategic decoupling and national control of ownership, a clear analogy to the early 20th century. In 2018, Australia financed the Coral Sea Cable to Papua New Guinea and the Solomon Islands, replacing a system already contracted to China's Huawei Marine (Australian Strategic Policy Institute, 2024). In 2021, US security objections led to the cancellation of a World Bank tender that China's HMN Tech had been positioned to win; the United States, Japan, and Australia then committed to finance the East Micronesia Cable themselves, with supply awarded to Japan's NEC (Jamestown Foundation, 2025). This policy approach was formalized in 2024, when seventeen governments, including the United States, the European Union, and Japan, endorsed a joint statement committing to "secure and verifiable" cable suppliers and to the development of alternative routes that increase the resilience of their networks (US Department of State, 2024).

These commitments are complemented by entry barriers: the United States now presumptively denies landing licenses to foreign adversary controlled applicants and prohibits Chinese equipment in new cable systems (Federal Communications Commission, 2025), with similar de facto bans by Australia and Japan. China pursues the same strategy from the other side. It has withheld permits for cable work in waters it claims in the South China Sea, delaying the SJC2 cable connecting Singapore to Japan by several years (Murphy and Bryja, 2025). Its PEACE cable, owned by a subsidiary of China's Hengtong Group, connects Asia to Europe via Africa on a route that bypasses India (Bloomberg News, 2025). And after US pressure shifted the SeaMeWe-6 supply contract from HMN Tech to the American firm SubCom, China's three state owned carriers withdrew from the consortium and began a rival system, the Europe-Middle East-Asia cable, supplied by HMN Tech with state subsidies (Brock, 2023).

7 Conclusion

The undersea cable infrastructure is the backbone of the global economy. This paper provides the first quantitative analysis of vulnerabilities and strategic dependencies in this global telecommunications network.

We assemble *GlobeTel*, a new dataset covering the universe of international submarine cables laid between 1850 and 2025, and embed it in a traffic-assignment model adapted from transportation economics. Our framework quantifies exposure to cable-related geopolitical risks such as embargoes or sanctions (dependence) and sabotage (vulnerability), identifying "digital chokepoints" and estimating potential welfare losses at the global and local level.

Three lessons emerge. First, the network has always been an instrument of power: hegemons have dominated it from British telegraphy in the 19th century to the US-centered internet today, and have repeatedly weaponized their control through espionage, exclusion, and sabotage. Second, the modern network remains highly vulnerable to large-scale disruption: cutting fifteen cables in a single corridor could sever links carrying up to a quarter of global data flows, with island economies most exposed. Third, states and firms respond: geopolitical tension has driven waves of defensive cable building for over a century, and is doing so again today. The stakes of these findings are rising, as AI and cloud computing deepen the world's dependence on bandwidth at the same time as underwater drones make the infrastructure that carries it cheaper to attack. We conclude that connectivity and vulnerability are two sides of the same coin. The gains from digital integration documented in the literature rest on physical infrastructure that is concentrated in a few corridors and controlled by a few firms and countries.

Our approach extends to other networks as well: because the dependence and vulnerability measures require only network shares and rerouting costs, the same framework can be used to identify and price chokepoints in pipelines, payment systems, shipping lanes, and any other infrastructure through which interdependence can be weaponized.

References

Ahvenainen, Jorma, *The Far Eastern Telegraphs: the History of Telegraphic Communications between the Far East, Europe, and America before the First World War*, Helsinki: Suomalainen Tiedeakatemia, 1981.

— , *The History of the Caribbean Telegraphs before the First World War*, Helsinki: Suomalainen Tiedeakatemia, 1996.

— , *The European Cable Companies in South America: Before the First World War*, Helsinki: Academia Scientiarum Fennica : Distributed by Tiedekirja, 2004. OCLC: ocm54512593.

— , *The History of the Near Eastern Telegraphs before the First World War*, Helsinki: Finnish Academy of Science and Letters, 2011. OCLC: ocn733408869.

Aldrich, Richard J., *GCHQ: The Uncensored Story of Britain's Most Secret Intelligence Agency*, London: HarperCollins Publishers, 2011.

Allen, James, "*Report of Col. James Allen, Chief Signal Officer, Army of Porto Rico, of the Operations of the Signal Corps on the South Coast of Cuba and in Porto Rico*," Report dated Ponce, Porto Rico, September 1, 1898 1898. The linked Atlantic Cable page reproduces an extract from the 1899 report to the Chief Signal Officer. Accessed 2026-03-19.

Allen, Treb and Costas Arkolakis, "Trade and the Topography of the Spatial Economy," *Quarterly Journal of Economics*, August 2014, *129* (3), 1085–1140.

— **and** — , "The Welfare Effects of Transportation Infrastructure Improvements," *The Review of Economic Studies*, 2022, *89* (6), 2911–2957.

Andrew, Christopher, *The Defence of the Realm: The Authorized History of MI5*, London New York, New York Toronto Dublin Camberwell New Delhi Rosedale, Auckland Johannesburg: Penguin, 2010.

Arkolakis, Costas and Conor Walsh, "Clean Growth," NBER Working Paper 31615 2023.

Australian Strategic Policy Institute, "Australians know about 'Fibre to the Premises', but what about 'Fibre to the Pacific'?," The Strategist, Australian Strategic Policy Institute 2024.

Bamford, James, *Body of Secrets: Anatomy of the Ultra-Secret National Security Agency*, New York: Anchor Books, 2002.

Begenau, Juliane et al., "Big Data in Finance and the Growth of Large Firms," *Journal of Monetary Economics*, 2018, *97*, 71–87.

Bertsch, Gary K., ed., *Controlling East-West Trade and Technology Transfer: Power, Politics, and Policies*, Durham, NC: Duke University Press Books, 1988.

Bhutani, Ankit et al., "The Missing Value of Data," in "NBER Macroeconomics Annual 2026, Volume 41," University of Chicago Press, 2026, chapter 1.

Bloomberg News, "Inside the US-China battle over undersea internet cables powering AI's future," Bloomberg 2025.

Bragadin, Marc' Antonio, *The Italian Navy in World War II*, Naval Institute Press, 1957. Internet Archive edition, accessed 2026-03-19.

Brock, Joe, "China plans $500 million subsea internet cable to rival US-backed project," Reuters, April 6 2023.

Broner, Fernando et al., "Hegemonic Globalization," 2025. CEPR Discussion Paper 20339.

Bruton, Elizabeth, "The Cable Wars: Military and State Surveillance of the British Telegraph Cable Network During World War One," *Marklund, Andreas and Mogens Rüdiger (ed.s). Historicizing Infrastructure (Aalborg: Aalborg University Press, 2017)*, January 2017.

Brynjolfsson, Erik, Avinash Collis, W. Erwin Diewert, Felix Eggers, and Kevin J. Fox, "GDP-B: Accounting for the Value of New and Free Goods," *American Economic Journal: Macroeconomics*, 2025, *17* (4), 312–344.

Caoui, El Hadi and Andrew Steck, "Diversification, Market Entry, and the Global Internet Backbone," *Review of Economic Studies*, 2025. Forthcoming. DOI: 10.1093/restud/rdaf097.

Carnegie Endowment, "Securing Europe's subsea data cables," Technical Report, Carnegie Endowment for International Peace, Washington, DC December 2024.

Clayton, Christopher, Antonio Coppola, Matteo Maggiori, and Jesse Schreger, "Geoeconomic Pressure," 2025. NBER Working Paper No. 34020.

—, Matteo Maggiori, and Jesse Schreger, "Putting Economics Back into Geoeconomics," in John V. Leahy and Valerie A. Ramey, eds., *NBER Macroeconomics Annual 2025, Volume 40*, Macroeconomics Annual, University of Chicago Press, 2025, chapter 2.

—, —, and —, "A Framework for Geoeconomics," *Econometrica*, 2026, *94* (1), 105–136.

Conley, Heather A., Sophie Arts, Kristine Berzina, and Frida Rintakumpu, "*Protecting Undersea Infrastructure in the North American Arctic: Lessons from Incidents in the Baltic Sea and High North*," May 2024.

Craven, John P., *The Silent War: The Cold War Battle Beneath the Sea*, 1st touchstone ed ed., New York, NY: Touchstone, 2002. OCLC: 49503087.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced trade," *American Economic Review*, 2007, *97* (2), 351–355.

Department of Veterans' Affairs, *Snapshots: Australians in War and Conflict*, Canberra: Department of Veterans' Affairs, September 2024.

Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, April 2018, *108* (4-5), 899–934.

— and Richard Hornbeck, "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, May 2016, *131* (2), 799–858.

Egorov, Konstantin et al., "Trade Sanctions," 2025. CEPR Discussion Paper 20601.

European Commission, "European Commission Submarine Cable Infrastructures Expert Group Report – Security and Resilience of EU Submarine Cable Infrastructures," Technical Report, European Commission, Brussels February 2026.

Fajgelbaum, Pablo D. and Edouard Schaal, "Optimal Transport Networks in Spatial Equilibrium," *Econometrica*, 2020, *88* (4), 1411–1452.

Farboodi, Maryam and Laura Veldkamp, "Long-Run Growth of Financial Data Technology," *American Economic Review*, 2020, *110* (8), 2485–2523.

Farrell, Henry and Abraham L. Newman, "Weaponized Interdependence: How Global Economic Networks Shape State Coercion," *International Security*, 2019, *44* (1), 42–79.

— and Abraham Newman, *Underground Empire: How America Weaponized the World Economy*, London: Allen Lane, 2023.

Federal Communications Commission, "Review of Submarine Cable Landing License Rules and Procedures To Assess Evolving National Security, Law Enforcement, Foreign Policy, and Trade Policy Risks," Technical Report FCC 25-49; 90 FR, October 27, 2025, Federal Communications Commission October 2025.

Felbermayr, Gabriel et al., "The Global Sanctions Data Base," *European Economic Review*, 2020, *129*, 103561.

Gallea, Quentin, Massimo Morelli, and Dominic Rohner, "Power in the Pipeline," mimeo October 2022.

Giles, Keir, "Russia's 'New' Tools for Confronting the West: Continuity and Innovation in Moscow's Exercise of Power | Chatham House – International Affairs Think Tank," 2016.

Goldfarb, Avi and Catherine Tucker, "Digital Economics," *Journal of Economic Literature*, 2019, *57* (1), 3–43.

Goodrich, Caspar F., "*The St. Louis' Cable-Cutting*," *U.S. Naval Institute Proceedings*, January 1900, *26* (1), 157–166. Vol. 26/1/93.

— , "*Cable Cutting*," *U.S. Naval Institute Proceedings*, January 1904, *30* (1).

Headrick, Daniel R., *The Tools of Empire: Technology and European Imperialism in the Nineteenth Century*, New York: Oxford University Press, 1981.

— , *The Invisible Weapon: Telecommunications and International Politics, 1851–1945*, New York: Oxford University Press, 1991.

— , *The Invisible Weapon: Telecommunications and International Politics, 1851-1945*, OUP USA, 1991. Google-Books-ID: UDXnCwAAQBAJ.

Hinsley, F. H., *British Intelligence in the Second World War: Its Influence on Strategy and Operations*, New York: Cambridge University Press, 1979.

Hjort, Jonas, "The Economic Impact of Internet Connectivity in Developing Countries," *Annual Review of Economics*, 2025. Forthcoming.

— **and Jonas Poulsen**, "The Arrival of Fast Internet and Employment in Africa," *American Economic Review*, 2019, *109* (3), 1032–1079.

— **and Lin Tian**, "The Economic Impact of Internet Connectivity in Developing Countries," *Annual Review of Economics*, August 2025, *17* (1), 99–124.

Hochfelder, David, *The Telegraph in America, 1832–1920*, Baltimore: Johns Hopkins University Press, 2013.

Hsu, Jason, "Testimony before the U.S.-China Economic and Security Review Commission," U.S.-China Economic and Security Review Commission March 2026. Starlink talks broke down over local majority-ownership rule; Musk acknowledged Beijing told him not to launch in China; jamming over a Taiwan-sized area judged feasible.

IEEE, "Meta's undersea cable project skirts conflict zones," *IEEE Spectrum*, March 2025.

Imbs, Jean and Laurent Pauwels, "Mapping the Global Submarine Cable Network: Resilience, Vulnerability, and Geopolitics," 2026. Working paper, NYUAD.

International Cable Protection Committee, "Cableships of the World," `https://www.iscpc.org/information/cableships-of-the-world/` 2025. Accessed March 2026.

Itskhoki, Oleg and Elina Ribakova, "The Economics of Sanctions: From Theory into Practice," *Brookings Papers on Economic Activity*, 2024, *55* (2), 425–470.

Jamestown Foundation, "Laying down the law under the sea: Analyzing the US and Chinese submarine cable governance regimes," China Brief, Jamestown Foundation 2025.

Jeon, Jihye and Marc Rysman, "Investment and Usage of the Subsea Internet Cable Network," 2025. Manuscript, Boston University.

Jones, Charles I. and Christopher Tonetti, "Nonrivalry and the Economics of Data," *American Economic Review*, 2020, *110* (9), 2819–2858.

Juhász, Réka and Claudia Steinwender, "Spinning the Web: The Impact of ICT on Trade in Intermediates and Technology Diffusion," Technical Report w24590, National Bureau of Economic Research, Cambridge, MA May 2018.

Kennedy, P. M., "Imperial Cable Communications and Strategy, 1870–1914 1," in "The War Plans of the Great Powers (RLE The First World War)," Routledge, 2014. Num Pages: 24.

Kimball, William W., "*Plan of Operations Against Spain*," 1896.

Kunert, A., *Telegraphen-Seekabel*, Vol. 2 of *Geschichte der deutschen Fernmeldekabel*, Deutsch-Atlantische Telegraphengesellschaft, Koeln, and Norddeutsche Seekabelwerke A.-G., Nordenham, 1962.

Mauldin, Alan, "The Red Sea: A key subsea cable crossroads under siege," TeleGeography 2024.

Mayer, Thierry et al., "The Fragmentation Paradox: De-risking Trade and Global Safety," 2025. CEPR Discussion Paper 20564.

McKinley, William, "Special Message to the Congress on the Need for Cable Communications with the Philippine Islands," in Gerhard Peters and John T. Woolley (eds.), *The American Presidency Project* February 1899. `https://www.presidency.ucsb.edu/`.

Middleton, Alexandra and Bjørn Rønning, "Geopolitics of subsea cables in the Arctic," *The Arctic Institute – Center for Circumpolar Security Studies*, 2022. August 2, 2022.

MODA, "Ministry of Digital Affairs (MODA), Resilience Construction Department, Mid-Term Infrastructure Case Plan," Review No. 112-3001-09-20-04, Executive Yuan, Taiwan 2023. Cable system transmission capacity ∼83 Tbps; satellite ∼2 Gbps, characterised as ≈1/41,500 of cable capacity.

—, "Non-Geostationary Satellite Orbit Emergency Response Network for Major Disasters, Now Available Nationwide," MODA Press Release 2024. 773 non-geostationary ground stations using OneWeb (LEO) and SES (MEO) for emergency government command communications.

Mohr, Cathrin and Christoph Trebesch, "Geoeconomics," *Annual Review of Economics*, 2025, *17* (1), 563–587.

Murphy, Erin L. and Matt Pearl, "China's Underwater Power Play: The PRC's New Subsea Cable-Cutting Ship Spooks International Security Experts," April 2025.

— **and Thomas Bryja**, "The strategic future of subsea cables: Singapore case study," Technical Report, Center for Strategic and International Studies, Washington, DC November 2025.

Naval Association of Australia, "*Monthly Ceremony: Jack's Memorial Southbank*," Event handout / commemorative pamphlet November 2020.

Nigmatulina, Dzhamilya, "Sanctions and Misallocation: How Sanctioned Firms Won and Russia Lost," 2022. CEP Discussion Paper 1886.

Ocon, Jaime and Jonathan Walberg, "*China's Undersea Cable Sabotage and Taiwan's Digital Vulnerabilities*," *Global Taiwan Brief*, 2025, *10* (11). Published by the Global Taiwan Institute; accessed 2026-03-19.

O'Connell, John F., "*German Special Submarine Operations During WW I*," *The Submarine Review*, April 2010. Naval Submarine League archive; accessed 2026-03-19.

Ottaviano, Gianmarco I. P., "Digital Network Centrality and the Structure of Goods Trade," CEP Discussion Paper dp2158, Centre for Economic Performance, LSE 2026.

Peters, Benjamin, *How Not to Network a Nation: The Uneasy History of the Soviet Internet*, The MIT Press, March 2016.

Royal Navy, "*Admiralty's War Diary World War Two*," n.d. Collection page for the Second World War Admiralty War Diaries.

Sontag, Sherry and Christopher Drew, *Blind Man's Bluff: The Untold Story of American Submarine Espionage*, New York: HarperTorch, 1999.

Starosielski, Nicole, *The Undersea Network*, Duke University Press, 2015.

Steinwender, Claudia, "Real Effects of Information Frictions: When the States and the Kingdom Became United," *American Economic Review*, March 2018, *108* (3), 657–696.

Subramanian, Samanth, *The Web Beneath the Waves: The Fragile Cables that Connect Our World*, Columbia Global Reports, 2025.

Sutherland, Ewan, "*The Yi Peng 3 and Eagle S incidents - cutting cables in the Baltic Sea*," in "33rd European Regional ITS Conference, Edinburgh, 2025: Digital innovation and transformation in uncertain times" International Telecommunications Society (ITS) Edinburgh 2025. Conference paper.

Taipei Times, "Taiwan-Starlink service talks fall through over regulatory issues," May 2026.

SpaceX cited >99% 4G/5G coverage as reason to deprioritize Taiwan; NCC talks failed over wholly owned entity demand.

The Economist, "The AI boom and geopolitics are rewiring Asia's oceans," *The Economist*, 2026. June 29, 2026.

TWNIC, "Taiwan Network Information Center (TWNIC), Internet Connection Bandwidth Survey, Q2 2025," July 2025. Taiwan's external connectivity carried by undersea cables; more than 99 percent of external traffic.

United States. Navy Department. Bureau of Navigation, *Appendix to the Report of the Chief of the Bureau of Navigation: Naval Operations of the War with Spain*, Washington, D.C.: United States Government Printing Office, 1898. Digitized by The Portal to Texas History.

US Department of State, "Joint statement on the security and resilience of undersea cables in a globally digitalized world," Office of the Spokesperson, US Department of State September 2024.

Veldkamp, Laura and Cindy Chung, "Data and the Aggregate Economy," *Journal of Economic Literature*, 2024, *62* (2), 458–484.

Wheeler, Tom, *Mr. Lincoln's T-Mails: The Untold Story of How Abraham Lincoln Used the Telegraph to Win the Civil War*, New York, N.Y: Harper Business, 2006.

Wilson, Herbert Wrigley, *The Downfall of Spain: Naval History of the Spanish-American War*, London: Sampson Low, Marston & Company, Limited, 1900. The linked page reproduces material from this book on the History of the Atlantic Cable & Undersea Communications site. Accessed 2026-03-19.

Winkler, Jonathan Reed, *Nexus: Strategic Communications and American Security in World War I*, Harvard University Press, 2008.

Wroblewski, Martin, *Moralische Eroberungen als Instrumente der Diplomatie: Die Informations- und Pressepolitik des Auswärtigen Amts 1902–1914*, Göttingen: V&R Unipress, 2016.

Online Appendix

Contents

Appendix A Digital blockades: historical narratives

The physical characteristics of subsea cables make them vulnerable not only to accidental damage but also to sanctions, sabotage, and espionage. This appendix documents the historical record of "digital blockades" and cable-based espionage linked to geopolitical tension, including periods of war. To compile this event archive we build on the *GlobeTel* dataset together with cable-specific documentation in our primary sources, such as the cable nomenclatures. Important general sources include Kunert (1962), Headrick (1981, 1991b), Sontag and Drew (1999), Winkler (2008), Kennedy (2014), Starosielski (2015), Giles (2016), and Bruton (2017).

The appendix is organized into three parts, corresponding to the three main channels through which cable control has been exercised: Section A.1 covers sanctions and the exclusion of geopolitical adversaries; Section A.2 covers espionage in peacetime and war; and Section A.3 covers cable cutting and sabotage.

A.1 Sanctions and exclusion of geopolitical adversaries

- *Telegraphy capture during the American Civil War (domestic cables):* During the American Civil War, the Union swiftly took control of much of the nation's telegraph network, effectively limiting the Confederacy's access. Prior to the war, telegraph lines were operated by private companies like the American Telegraph Company and Western Union, facilitating domestic and some limited international communication. In 1861, the Union established the US Military Telegraph Corps, centralizing control and curtailing the South's ability to coordinate military operations. By seizing key telegraph hubs, especially around Washington D.C., the Union secured a significant strategic advantage (Hochfelder, 2013; Wheeler, 2006).

- *The weaponization of telegraphy pre-WW1 (colonial wars):* In the late 19th century and the run-up to WW1, Great Britain increasingly exploited its dominance over global telegraphy to secure military and strategic advantages, prompting frustration and diplomatic protest from rival powers.
  - An illustrative case is the *Fashoda Incident* of 1898, when Britain and France confronted each other over territorial claims in the Upper Nile region (present-day South Sudan). British forces maintained direct telegraph access via a line running from Omdurman through Khartoum and Cairo to London, allowing for rapid military communication. In

contrast, the French lacked their own telegraph connection, relying instead on slow couriers. During the incident, the British denied the French access to the telegraph line, resulting in an important strategic disadvantage. According to Headrick (1991), British commander General Kitchener transmitted selectively negative information about the French position, leading to a misperception in Paris that Marchand's expedition was in serious difficulty. This contributed to their decision to back down.

  - A far more consequential cable blockade occurred during the *Boer War* (1899–1902), when Britain fought against the Boer republics of Transvaal and the Orange Free State. Beginning in October 1899, British authorities implemented strict censorship, re-routing, and outright blocking of telegraph traffic to and from southern Africa. The primary aim was to suppress independent war reporting and disrupt diplomatic communications between the Boers and their European sympathizers. The ban was so broad that it prohibited the transmission of any ciphered or foreign-language telegram south of Aden, forcing such messages to pass through London for approval or redaction. As a result, the cable blockade also curtailed French and German communication with their own African colonies, prompting formal diplomatic protests and ongoing criticism from European governments.
  - *Britain's offensive cable strategy pre-WW1:* Following Fashoda and the Boer War, British strategic thinking increasingly evolved towards actively exploiting its dominance over global telecommunications. This shift was facilitated by the completion of the "All Red Line" in 1902, an extensive global network of submarine cables connecting British territories without relying on foreign cables or stations. Britain's intent to weaponize its cable dominance was codified in key policy documents, including the 1904 *Memorandum on the Protection of British Submarine Cables and on the Destruction of the Enemy's Cables*, which outlined plans to disrupt adversaries' cable networks. The 1911 report *Submarine Cable Communications in Times of War* further detailed Britain's options to control communications, intercept enemy messages, isolate rival powers, and build resilience.

- *Soviet exclusion during the Cold War:* During the Cold War, the US and the UK effectively excluded the Soviet Union and its allies from accessing the growing, Western-controlled global submarine cable network, as well as the INTELSAT satellite system. While not subject to a formal telecommunications embargo, the Soviet bloc could therefore not engage in global telecommunication via Western routes. In addition, the export of advanced telecommunications technologies was tightly controlled via mechanisms such as the Coordinating Committee for Multilateral Export Controls (CoCom). As a result, the USSR had to rely on intra-bloc terrestrial cables, radio communications, and its own satellite systems such as Molniya and Gorizont, which were similarly off-limits to Western nations (Peters, 2016; Bertsch, ed, 1988).

- *Digital embargo of Cuba:* Cuba has long faced a telecommunications and cable embargo, severely restricting its telephony and internet access. After the 1959 revolution, the US banned cable links with Cuba and later imposed strict export controls on telecom technologies, along with more recent restrictions on internet services such as Google Earth, Zoom, and Amazon. Our data show that all US-owned cables in the Gulf of Mexico and the Caribbean bypass Cuba. Consequently, until the mid-2000s, Cuba could access the internet solely via geostationary satellites, at high cost and limited bandwidth. It was not until 2013 that Cuba connected to a subsea fiber-optic cable via Venezuela, followed by a second link to Martinique in 2023. Although the US lifted its cable-laying ban in 2016, Cuba's 2022 request to access the US-owned ARCOS-1 system was denied by the US authorities on national security grounds.

A.2 Espionage in peacetime and war

- *British wiretapping in WW1 (& the Zimmermann Telegram):* During WW1, British intelligence systematically tapped undersea telegraph cables by accessing landing sites like Porthcurno and intercepting messages at relay stations in London or Cornwall. After the UK cut all of Germany's direct cables to the Americas (see below), German messages could only be transmitted via neutral countries, enabling interception at British-controlled hubs, or by high-powered radio, which could be deciphered. This decryption was handled by Room 40, a secret Admiralty unit specializing in naval and diplomatic codebreaking. A major success was the interception of the Zimmermann Telegram, in which Germany proposed a military alliance with Mexico against the US. The interception and decryption of this telegram helped shift American opinion toward entering the war against Germany (Hinsley, 1979).

- *British and US communication espionage in WW2:* During WW2, Allied intelligence services, primarily Britain's Bletchley Park and the US Signal Intelligence Service, intercepted both high-frequency radio and undersea cable traffic used by the Axis powers. These channels carried encrypted communications on military logistics, troop movements, and diplomacy. A major breakthrough came with the decryption of German Enigma traffic under the ULTRA program, which significantly aided the Allies (Andrew, 2010).

- *Cold War espionage I — Operation Ivy Bells (1971–1981):* In 1971, operatives of the US Navy and National Security Agency (NSA) covertly tapped a Soviet undersea communications cable in the Sea of Okhotsk. The cable transmitted sensitive information between the Soviet Pacific Fleet headquarters in Vladivostok and naval installations on the Kamchatka Peninsula. The operation yielded high-value intelligence on Soviet naval strategy and submarine deployments for nearly a decade, until it was compromised by former NSA contractor Ronald Pelton in 1980 (Craven, 2002; Sontag and Drew, 1999).

- *Cold War espionage II — the ECHELON program (Five Eyes):* Established in the 1970s, the ECHELON program, operated by the Five Eyes alliance (USA, UK, Canada, Australia, New Zealand), was designed to intercept global communications for signals intelligence, initially targeting the Soviet bloc. The system primarily intercepted satellite transmissions, but also gathered data from international cable networks where signals were routed through key landing points. Over time, ECHELON expanded its collection to include a wide range of diplomatic, commercial, and military communications worldwide (Bamford, 2002; Aldrich, 2011).

- *The NSA Tempora program tapping internet cables (ca. 2009–2013):* Revealed by Edward Snowden in 2013, the NSA and the UK's GCHQ tapped into fiber-optic subsea cables at landing stations to collect vast amounts of internet and telephone data. By accessing cable data streams directly, these agencies gained unprecedented surveillance capabilities. The case underscores the scale and technical sophistication of modern cable-based espionage. Other key NSA programs involving international subsea cable interception include MUSCULAR, which tapped private fiber links between major data centers operated by Google and Yahoo, and Fairview, which targeted international communications carried through US infrastructure, including subsea cable landing stations.

A.3 Cable cutting and sabotage

By systematically searching archives, scientific literature, and online sources for historical cable cutting cases, we identified at least 91 incidents of attempted or successful subsea cable cutting up to 1945, which should be seen as a lower bound. We find that until the end of WW2, reported cases of subsea cable cutting appeared almost exclusively as outright acts of war during major conflicts. This changed during the Cold War, when the US and the Soviet Union prioritized espionage over outright sabotage and the boundaries between deliberate aggression and accidental damage became increasingly blurred by "gray-zone" attacks.

- *Sabotage in the Spanish-American War:* One of the earliest conflicts in which cable cutting became a central strategic instrument was the Spanish-American War of 1898. Well before hostilities began, US Navy planners had recognized the military importance of the telegraph cables connecting Spain to its colonies in Central America (Kimball, 1896). Once the conflict started, the US Navy systematically severed Spanish telegraph lines around Cuba and the Caribbean. Our data record at least 19 attempted or successful cable cuts. These operations disconnected the Spanish colonies of Cuba and Puerto Rico from Madrid, depriving Spain of command and control over its forces and giving the United States a decisive advantage.[11]

- *Sabotage in WW1:* In August 1914, the Royal Navy successfully severed the main telegraph cables that linked Germany to the rest of the world, in particular those running through the English Channel to the US, Brest, Vigo, the Azores, and Tenerife (Kunert, 1962).[12] For the remainder of the war, Germany had only limited means to communicate or transfer funds internationally, including with its overseas territories and allies, and to counter the intense propaganda that the UK transmitted to the US.

[11] Insightful sources on this episode include United States. Navy Department. Bureau of Navigation (1898), Goodrich (1900), Goodrich (1904), Wilson (1900), and Allen (1898).

[12] These operations were based on prewar plans in which the Royal Navy identified the most important German telegraph lines to be cut. Great Britain also occupied or destroyed telegraph stations such as the one on Yap Island, which was crucial for Germany's Pacific communication (Winkler, 2008).

Germany retaliated by cutting the Baltic Sea cables that connected Britain and France with Russia (Winkler, 2008). Moreover, towards the end of the war, Germany began to deploy submarines to sever cables and threaten repair ships (O'Connell, 2010). This strategy proved increasingly effective: the Atlantic became a telegraphy bottleneck in which commercial, diplomatic, and military traffic competed for dwindling capacity. According to historians, this experience prompted the United States to invest heavily in building its own international cable infrastructure and radio technology (Winkler, 2008).

- *Sabotage in WW2:* As in WW1, the belligerents of WW2 spent considerable resources to disrupt one another's submarine cable networks. A central objective was to force the enemy to rely on radio communication (wireless telegraphy), which was widely used in this period but much easier to intercept than cable communication. In 1939, the British cut German cables from Emden to Spain, Portugal, and the US (via the Azores), followed by those linking the German islands of Borkum and Norderney with England and Norway. During the first year of the war there were at least 25 incidents of cable cutting or attempted cutting of German cables in the English Channel, the North Sea, and the Atlantic. The British also launched several operations against Italy and Japan, in particular by cutting Italy's cables in the Atlantic and Mediterranean and severing multiple Japanese lines.[13] While Germany largely refrained from cable cutting, Italy cut all British cables around Malta and Gibraltar, and Japan captured and cut British cables in the Far East (Bragadin, 1957). Towards the end of the war in the Pacific, specially developed Australian midget submarines cut the remaining Japanese cables to Saigon, Singapore, and Hong Kong to enforce the use of radio communication ("Operation Foil" and "Operation Sabre") (Naval Association of Australia, 2020). These operations were militarily crucial, as they allowed the US to intercept Japan's strategic planning prior to dropping the atomic bombs in August 1945 (Department of Veterans' Affairs, 2024).

- *Sabotage in the Cold War:* After 1945, deliberate cable cutting between adversaries became less common, as major powers such as the US and the Soviet Union prioritized espionage over outright sabotage. A notable

[13] The most valuable sources for British WW2 cable cutting are the Royal Navy's War Diaries: Royal Navy (n.d.).

exception is the 1959 Novorossiysk incident, in which the US Navy boarded a Soviet fishing trawler suspected of having cut five transatlantic cables operated by AT&T and Western Union, roughly 120 miles off the coast of St. John's, Newfoundland.

- *Baltic and North Sea incidents since 2022:* In the wake of escalating tensions and Russia's full-scale invasion of Ukraine in 2022, both concerns over and actual cable cutting incidents in Northern Europe increased.[14] In January 2022, the fiber-optic Svalbard Cable, owned by Space Norway and connecting mainland Norway to the Arctic island of Svalbard, was severed, reportedly by Russian fishing vessels (Conley et al., 2024). Svalbard is strategically located and hosts the SvalSat antenna park, the world's largest satellite ground station, serving NASA, ESA, and others. In October 2023, the Chinese-flagged vessel NewNew Polar Bear reportedly damaged two subsea cables linking Estonia with Finland and Sweden (Conley et al., 2024). In November 2024, Finland's C-Lion1 cable to Germany and the BCS East-West cable between Sweden's Gotland Island and Lithuania were severed, reportedly by the Chinese-flagged Yi Peng 3 carrying Russian fertilizer (Sutherland, 2025). In December 2024, the Estlink 2 power cable connecting Finland and Estonia and four telecom cables were cut by the Cook Islands-flagged Eagle S, linked to Russia's "shadow fleet" (Sutherland, 2025). In December 2025, three additional cables were cut or damaged, reportedly by a Russian cargo ship. While these incidents all involved ships linked to China and Russia, there is insufficient proof that they were deliberate or politically ordered. They are, however, often seen as part of a broader strategy of "gray-zone" or "hybrid" warfare (Farrell and Newman, 2023).

- *Taiwan incidents in recent years:* Taiwan has long blamed China for repeatedly damaging the subsea cables connecting the island to the outside world. In the Matsu Island incident of February 2023, the only two cables to Matsu were severed, reportedly by a Chinese fishing vessel and a cargo ship, disrupting internet access for the island's 14,000 inhabitants for several weeks (Ocon and Walberg, 2025). In January 2025, the Shunxing 39, a Chinese-crewed cargo ship, was suspected of damaging a subsea cable linking Taiwan to the US (Ocon and Walberg, 2025). In February 2025, the Hongtai 58, also with a Chinese crew, was detained following damage to a

[14]For an overview of the debate and the known facts, see the collection of articles by the Carnegie Endowment: `https://carnegieendowment.org/the-geopolitics-of-subsea-data-cables?lang=en`.

cable connecting Taiwan to the Penghu Islands (Ocon and Walberg, 2025). As with the Baltic Sea incidents, proof of deliberate sabotage is difficult, but Taiwanese authorities consider it likely (Murphy and Pearl, 2025).

Appendix B Empirical Implementation of the Traffic-Assignment Model

B.1 Physical network: reconstructing the seabed path

Cable workbooks report landing points, ownership, capacity, vintage, and in many cases declared total system length, but not the path a cable traces on the ocean floor. We reconstruct that path by tracing each segment as a least-cost route over a cost surface built from GEBCO bathymetry, so that fitted routes respond to seabed relief the way engineered cables do, and we validate the reconstruction against declared system length.

Table B1. Bias and accuracy of predicted submarine system lengths

Statistic	Value
Systems in graph with observed length	604
Median predicted/observed length	1.032
Aggregate predicted/observed length	0.967
Median ratio, systems $\geq$500 km	0.999

Notes: Cable system as the unit (2020 traffic-assignment graph against declared lengths in the 2026-04-10 consensus workbook). The bias is two-sided—350 systems fit long, 254 short—and sub-percent in aggregate.

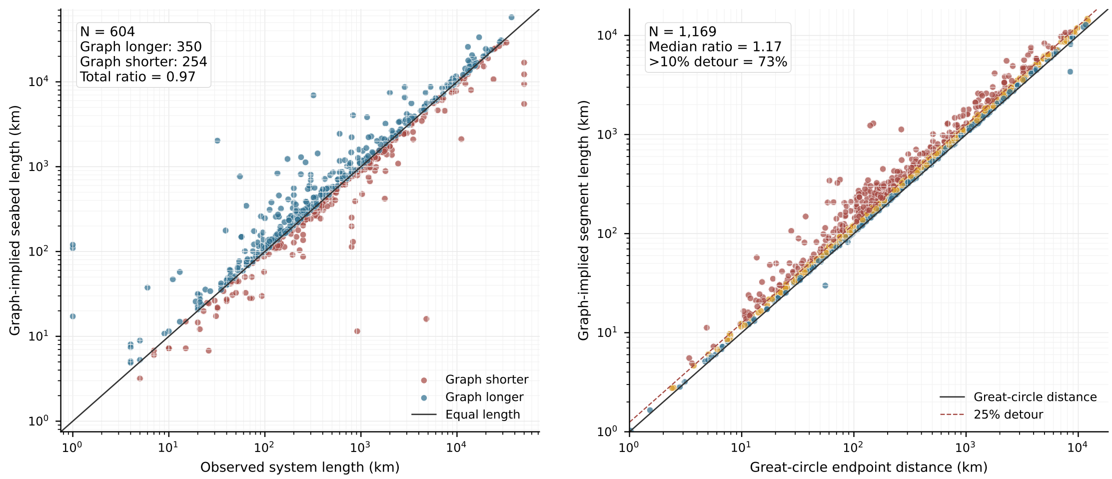

Figure B1. The fitted graph recovers real cable geometry. Left: graph-implied length tracks declared system length (604 systems, median ratio 1.03). Right: fitted segments run about 17% longer than the great circle at the median—the seabed detour.

B.2 Physical network: reconstructing the terrestrial network

For each candidate terrestrial configuration g we simulate the model globally, obtain the predicted free-flow round-trip time $\hat{\tau}_{njt}^{0,g}$ for every city pair and year, and compare it with the minimum observed round-trip time $\tau_{njt}^{\min}$, the low-congestion latency floor. The selected configuration minimizes the root-mean-squared error over the 18,329 matched country-pair-years,

$$\mathrm{RMSE}(g) = \left[\frac{1}{N_g} \sum_{(n,j,t)} \left(\tau_{njt}^{\min} - \hat{\tau}_{njt}^{0,g} \right)^2 \right]^{1/2}, \tag{B1}$$

which falls from 46.4 ms for a submarine-only network to 29.7 ms for the selected overlay ($R^2 = 0.90$).

Figure B2. Candidate terrestrial configurations by region, from national corridors to dense international meshes.

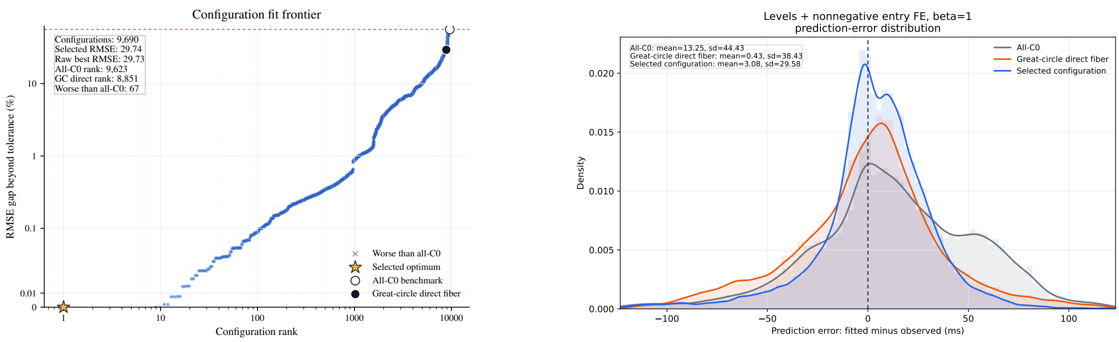

Figure B3. The selected overlay is the fit optimum. Left: latency error by configuration rank across 9,690 candidates; the national and great-circle benchmarks rank near the bottom. Right: the selected configuration's prediction error is centered on zero.

Appendix C Estimating the Model and Computing the Welfare Measures

C.1 Congestion

The volume–delay function turns an edge's load into delay,

$$\tau_\ell = \tau_\ell^0 \left[1 + \alpha \left(\frac{\text{Traffic}_\ell}{\text{Cap}_\ell} \right)^\beta \right], \tag{C1}$$

and the restriction that delay doubles at twice capacity, $2 = 1 + \alpha\, 2^\beta$, ties $\beta = -\log_2 \alpha$ and leaves α as the only free parameter. We recover it from the daylight cycle. For each origin–destination route p we push the hour-h demand profile through the assignment and form the congestion index $X_p(h;\beta) = \sum_{\ell \in p} \tau_\ell^0 \,(\text{Traffic}_\ell(h)/\text{Cap}_\ell)^\beta$, the excess delay the route's links carry at hour h, and regress observed latency on it route by route,

$$\tau_{p,h,d} = \gamma_{p,d} + \alpha\, X_p(h;\beta) + \varepsilon_{p,h,d}. \tag{C2}$$

The route-day effect $\gamma_{p,d}$ absorbs the free-flow level, the equipment, and the route's quality, so that within a route on a given day only the changing fullness of the shared links can move τ—the exogenous swing the daylight cycle supplies. Because X_p itself depends on β and on the equilibrium loads, α and the assignment are solved as a fixed point. The headline is the median of the route-level slopes in (C2), one vote per route so that a few high-leverage routes cannot drive it, taken over the within-capacity routes of the clean 2022–2023 panel with 90% intervals from a route-level bootstrap; Table C1 shows it holds across the alternative within-route estimators and year by year.

Table C1. Congestion coefficient across estimators and years. The curvature is implied by the doubling restriction $\alpha = 2^{-\beta}$; the estimating equation and sample are in the text.

Estimator	Congestion coefficient	Curvature (implied)
Within-route median slope (2022–2023)	0.25	2.0
Aggregate diurnal, OLS of medians (2022–2023)	0.22	2.2
Pooled within-route, multi-year (2018–2023)	0.24	2.1
Headline	**0.24**	**2.0**

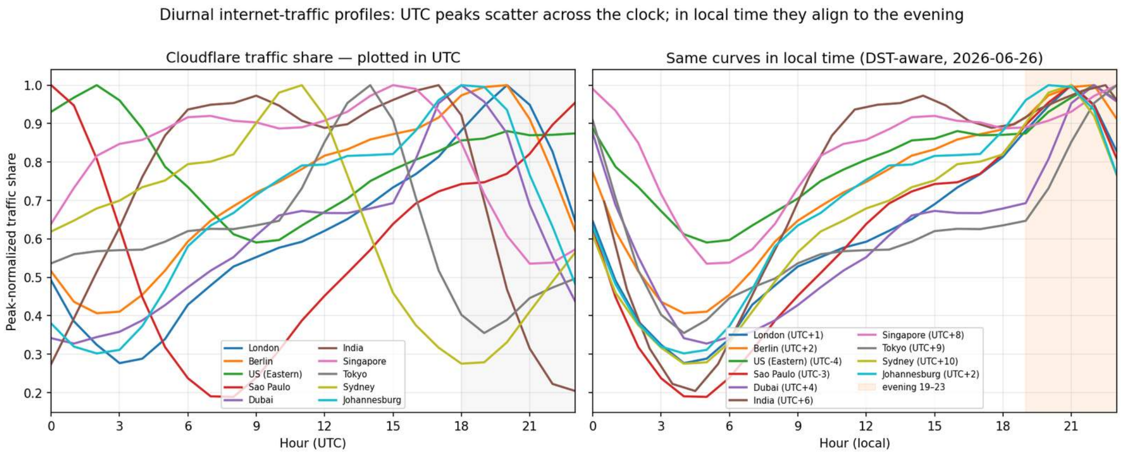

Figure C1. The daylight cycle as identifying variation. Left: peak-normalized traffic share by hour UTC—national busy hours scatter across the clock. Right: in local time they align to a common evening peak. This within-day rhythm supplies the congestion variation that identifies the coefficient.

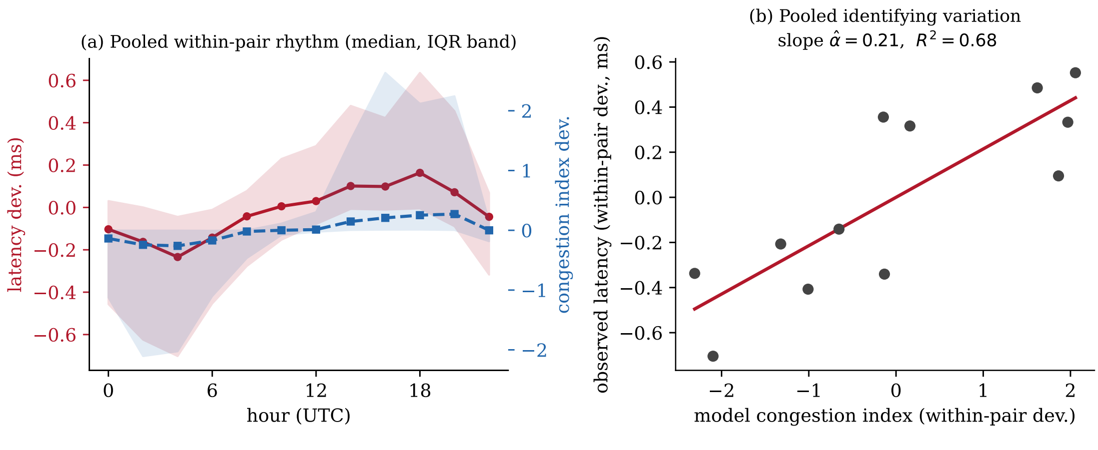

Figure C2. Pooled identification. Left: median within-route profile of observed latency (red) and the model congestion index (blue), interquartile bands; within-route correlations run 0.82–0.96. Right: within-route deviations give a slope of 0.21, and the median estimator delivers the headline 0.24.

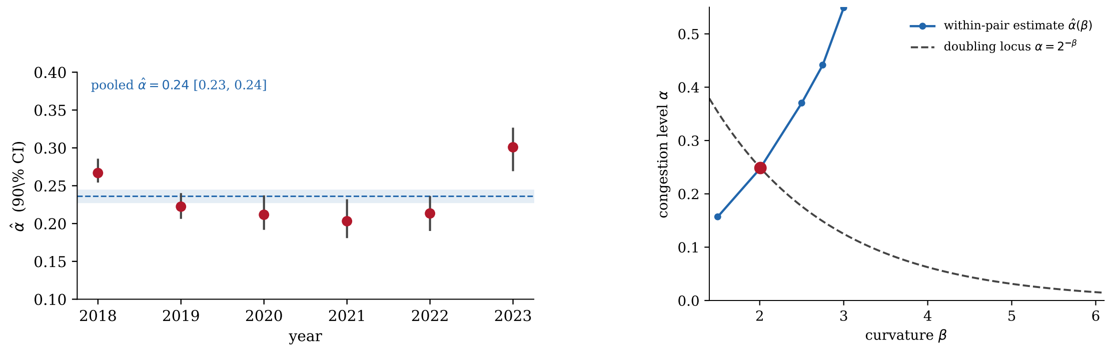

Figure C3. Stability and the doubling restriction. Left: year-by-year estimates around the pooled 0.24 (band), 2018–2023. Right: the within-route estimate meets the doubling locus at a curvature near 2.0.

C.2 Demand elasticities

As is standard in the gravity literature, we relate bilateral flows to the latency resistance term,

$$\log x_{njt} = \delta_{nt} + \delta_{jt} - \rho \log \tau_{njt} + \Gamma' Z_{njt} + \varepsilon_{njt}, \tag{C3}$$

where x_{njt} is the flow between n and j, τ_{njt} the observed round-trip time, Z_{njt} the usual bilateral covariates, and ρ the latency elasticity of demand. Observed latency is endogenous—congestion rises with the very flows on the left—so we instrument $\log \tau_{njt}$ with the free-flow latency $\log \tau^0_{njt}$, which is fixed by topology and fibre length and therefore shifts flows only through realized latency. We fit (C3) on the 455 undirected country pairs with positive 2023 traffic, by ordinary least squares and, in the instrumented columns, two-stage least squares, with heteroskedasticity-robust standard errors. The no-distance columns are the baseline: free-flow latency is nearly proportional to distance, so conditioning on distance strips out the very variation that identifies ρ.

Table C2. Latency elasticity of demand

	OLS	OLS	Instrumental variables	
	(1)	(2)	(3)	(4)
Latency elasticity of demand	1.720***	2.211***	2.453***	3.515***
	(0.191)	(0.233)	(0.254)	(0.402)
First-stage F-statistic			1702.1	178.1
Country-year effects	No	Yes	No	Yes
Observations	455	455	455	455

Notes: Undirected country pairs with positive 2023 traffic. All columns include bilateral controls; distance is excluded (see text). Columns (3)–(4) instrument observed round-trip time with free-flow latency. Heteroskedasticity-robust standard errors in parentheses. *, **, *** denote 10, 5, and 1 percent significance.

C.3 Prices and the output elasticity of data

The model prices transport in proportion to latency, a shape restriction we test by

$$\log p_{njt} = \log \varphi_t + b \log \tau_{njt} + u_{njt}, \qquad H_0\colon b = 1, \tag{C4}$$

where p_{njt} is the price per Gbps on route $n-j$ and φ_t the transport cost per unit of latency; this is a shape restriction, not a causal claim. The elasticity b is identified across routes within a capacity-and-year cell, and the sharpest form adds origin- and destination-city effects so that access and termination costs, which the model does not price into transport, drop out. We estimate (C4) on leased-line quotes at the route-by-capacity-by-quarter level (2012–2020; 2,258 routes in the 2020 cross-section) by ordinary least squares, and by two-stage least squares instrumenting realized latency with free-flow latency, with robust standard errors.

Once proportionality holds, prices identify the output elasticity of data, η, through how a unit of international access is priced. Three cases nest one another—competitive transport, a competitive (Bertrand) access fee, and a gatekeeper fee,

$$p_{nj} = \varphi\, \tau_{nj}, \qquad a_n^{\text{comp}} = \varphi\, \bar{\tau}_n, \qquad a_n^{\text{gate}} = \frac{\varphi}{\eta}\, \bar{\tau}_n, \tag{C5}$$

where a_n is country n's access price and $\bar{\tau}_n = \sum_j (x_{nj}/X_n)\, \tau_{nj}$ its traffic-weighted latency to international sources. Selling against data demand of elasticity $1/(1-\eta)$, the gatekeeper charges the inverse-elasticity markup over average cost, so the last case inverts to $\eta = \varphi\, \bar{\tau}_n / a_n$: the inverse gross markup of the access price over transport cost. All three share the same φ and $\bar{\tau}_n$; only the markup $1/\eta$ separates the gatekeeper from the competitive case, so comparing access prices across market structures, holding cost fixed, identifies η.

Imposing proportionality in (C4) recovers the transport cost as the median of $\log p - \log \tau$: about \$11.6 per Gbps·ms·month for a 10G circuit in 2020, down from about \$34 in 2013. We then read η off the gatekeeper markup, differencing it out of the access-to-latency ratio,

$$\log a_{nt} - \log \bar{\tau}_{nt} = c + \lambda\, G_{nt} + \delta_t + \varepsilon_{nt}, \qquad \hat{\eta} = e^{-\lambda}, \tag{C6}$$

where G_{nt} flags a concentrated market—an ownership Herfindahl of at least 0.5, built from cable-owner nationalities in the workbook the model uses for its dependence scenario, with the European Union pooled and a state incumbent where a country has no home firm. Estimated across 298 country-years (80 countries, 2020–2023) by ordinary least squares with year effects and standard errors clustered by country, this gives $\hat{\eta} = 0.58$ (0.35–0.75 across concentration thresholds). The design needs no cost level: by (C5) the access-to-latency ratio is the markup itself—1 in competitive markets, $1/\eta$ in gatekeeper ones—so differencing removes φ and any technology or utilization level common to both groups. Its one identifying assumption is that market structure is uncorrelated with that utilization wedge given the year. The natural concern is that concentrated markets are also small, so a per-unit fee could embed the fixed per-circuit charge spread over little volume; controlling for market volume, restricting to gatekeeper markets above the competitive-median size, and measuring concentration only where home firms exist all leave $\hat{\eta}$ at 0.52–0.69 (Table C3), so the gradient is market power, not fixed-cost amortization.

The markup recovers the model's η, but the accounting share usually reported is smaller, and the gap is informative. Write the generalized cost of a unit of access as

$$\kappa_{nj} = p_{nj} + w_n \chi\, \tau_{nj}, \tag{C7}$$

the money price plus the wage value $w_n \chi$ of the worker time a latency τ_{nj} ties up. Under the Cobb–Douglas technology η is the input's share of this generalized cost, $\eta = \sum_j \kappa_{nj} x_{nj} / Y_n$, whereas an expenditure share sees only the money part,

$$s_n^{\text{money}} = \frac{\sum_j p_{nj} x_{nj}}{Y_n} = \theta_n\, \eta, \qquad \theta_n \in (0,1). \tag{C8}$$

The money-expenditure rungs—telecom spending (about 1.3% of GDP), rising as one adds software, intangibles, then research and intellectual property—are therefore lower bounds on η, missing the time cost paid in hours rather than dollars; with $\hat{\eta} \approx 0.5$–0.6 and $s^{\text{money}} \approx 0.05$, money is under a tenth of the true cost of data. Because a smaller η implies a larger markup in (C5), any force that compresses the observed markup—regulation, large-buyer bypass—inflates $\hat{\eta} = e^{-\lambda}$, so the gatekeeper estimate brackets η from above while the expenditure rungs bracket it from below. With the coherence requirement that the aggregate access elasticity nest inside the cross-route trade elasticity, $\eta \leq 1 - 1/\rho = 0.59$ at $\rho = 2.45$, the feasible set is $[0.31, 0.59]$, and we adopt a headline welfare value of $\eta = 0.5$.

Table C3. Output elasticity of data, from gatekeeper access markups

Specification	Output elasticity of data	(Std. err.)	Observations
Baseline (year effects)	0.58	(0.19)	298
+ log GDP per capita	0.62	(0.18)	298
+ log market volume	0.69	(0.21)	294
+ market volume and GDP per capita	0.68	(0.20)	294
Gatekeeper markets above competitive-median size	0.65	(0.33)	201
Concentration only where home firms exist	0.52	(0.19)	282

Notes: Country-years, 2020–2023 (80 countries). Each row adds the noted control to the baseline; the elasticity is $e^{-\lambda}$ from the gatekeeper markup gap (see text). Standard errors clustered by country in parentheses.

Table C4. Price–latency elasticity: city cross-section and country panel

	City cross-section (2020)			Country panel (2012–2020)		
	Pooled	10G	100G	Year	Capacity×year	All products
	(1)	(2)	(3)	(4)	(5)	(6)
Latency elasticity	1.260	1.165	1.625	1.073	0.952	0.710
	(0.098)	(0.128)	(0.122)	(0.066)	(0.057)	(0.049)
p-value (elasticity = 1)	0.01	0.20	0.00	0.27	0.40	0.00
Capacity effects	Yes	–	–	–	Yes	Yes
Year effects	–	–	–	Yes	Yes	Yes
Latency measure	traceroute RTT			WonderProxy RTT		
Observations	2258	812	689	11511	11511	23026
R^2	0.59	0.34	0.45	0.48	0.73	0.83

Notes: Route by capacity product by quarter (a city pair). Log price per Gbps, ordinary least squares, robust standard errors. Proportionality is an elasticity of one (p-value row).

Table C5. Price–latency elasticity with city-endpoint effects

Specification	Latency elasticity	(Std. err.)	Observations
City-endpoint effects (baseline)	1.000	(0.084)	2258
+ log great-circle city distance	0.781	(0.151)	2258
(distance coefficient)	0.182	(0.101)	
Instrumented with free-flow latency (first-stage $F = 106$)	1.208	(0.101)	2254
Instrumented with speed-of-light minimum ($F = 182$)	1.109	(0.097)	2258

Notes: Routes (city pairs), 2020. Both endpoint-city effects absorbed, so access and termination costs drop out. Robust standard errors; proportionality (elasticity one) holds and survives instrumenting with free-flow latency.

C.4 Validation: traceroutes

The observed side is RIPE Atlas traceroutes mapped to submarine cables through the Nautilus crosswalk. We compare model-predicted and traceroute-implied cable use at the cable level and the pair-cable level, and check where satellites appear along a route.

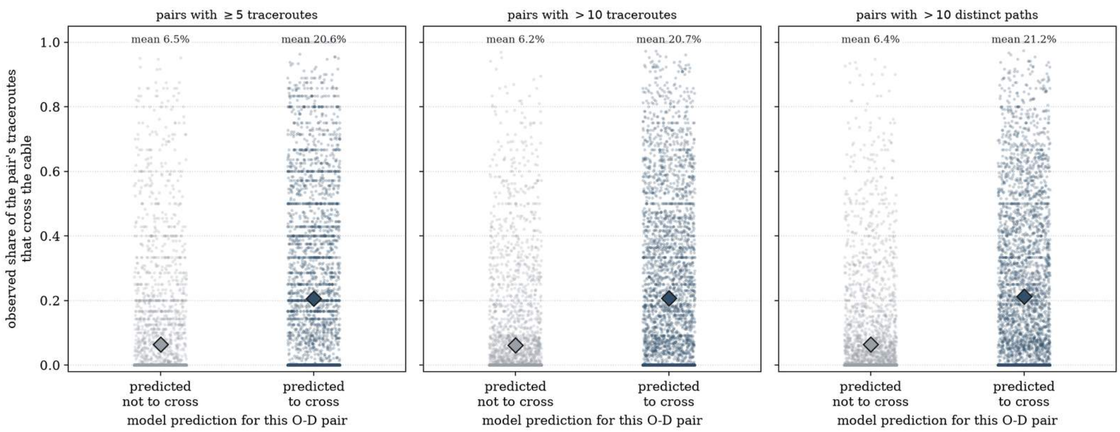

Figure C4. Predicted cables are observed far more often than unpredicted ones. Observed share of a pair's traceroutes crossing a cable, split by the model's prediction (diamonds mark means); panels add coverage thresholds.

C.5 Lit capacity

If the model routes traffic correctly, operators should light more of a cable's design capacity on the systems it loads more heavily. The informative test is on the utilization ratio,

$$\log \frac{L_{st}^{\text{lit}}}{L_{st}^{\text{pot}}} = b \log u_{st}^{\text{model}} + \delta_t + \varepsilon_{st}, \tag{C9}$$

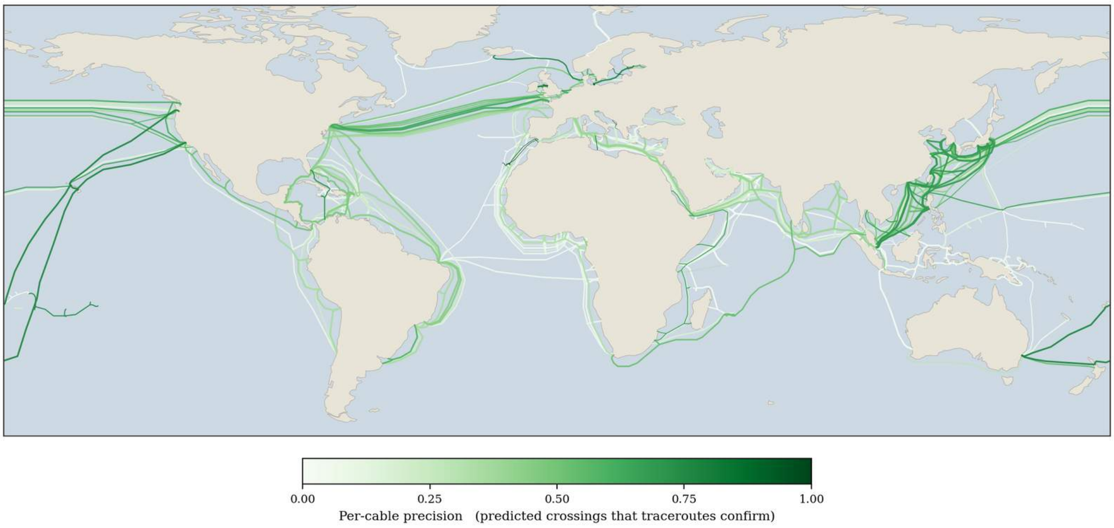

Figure C5. Per-cable routing precision: the share of each system's model-predicted city-pair crossings that traceroutes confirm.

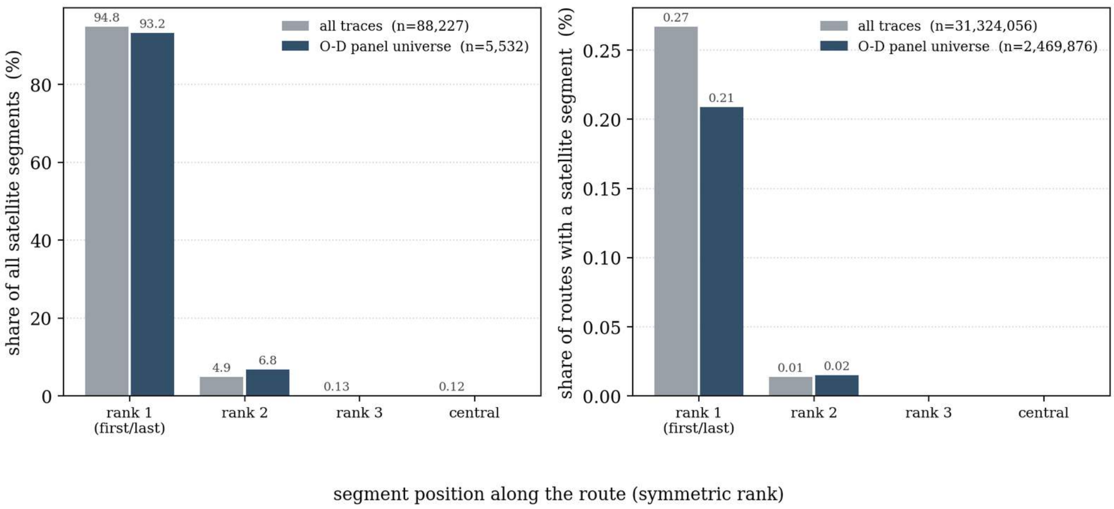

Figure C6. Satellites appear only as access legs. Left: 93–95% of observed satellite segments sit on the first or last hop, almost none mid-path. Right: only about 0.3% of routes carry any satellite segment.

Table C6. Cable-level routing validation against traceroute-implied usage shares

Comparison	Spearman correlation	Cables
Fractional demand-weighted share	0.636	172
Binary any-crossing share	0.610	172

Notes: Cable-level, RIPE Atlas / Nautilus traceroute-to-cable evidence (current demand-weighted 2022 run). Robustness variants—full-window weights, capital endpoints, unordered dyads—give 0.50–0.51.

Table C7. Pair-level routing precision versus a geography-aware random benchmark

Sample	Pairs	Actual mean	Actual pooled	Random expected	Normalized gain
All predicted pairs	270,102	34.5%	30.7%	6.5%	0.307
Pairs with ≥ 5 traceroutes	48,952	61.1%	54.0%	14.3%	0.549
Pairs with > 10 traceroutes	23,274	68.7%	61.3%	18.0%	0.618

Notes: Origin–destination pairs. Precision is the share of predicted mapped cables seen in at least one traceroute; the benchmark draws the same number of cables from the region-pair mapped-cable vocabulary.

where L^{lit} and L^{pot} are lit and potential (design) capacity and u_{st}^{model} the model's peak utilization of system s. Dividing both sides by design capacity strips out the mechanical channel through which large cables draw both more real and more model traffic—which is why the levels regression, shown for contrast, loses its slope once design capacity is controlled. We fit (C9) across 331 system-years (46 harmonized systems, 2014–2022) by ordinary least squares with year effects, clustering standard errors by system.

Table C8. Lit-capacity validation, yearly panel 2014–2022

Specification	Systems	Slope	Std. err.	R^2	Spearman
Lit-to-potential ratio on model utilization (logs)	331	0.143	0.071	0.126	0.355
credible-utilization subsample	253	0.088	0.058	0.146	0.285
Levels: lit capacity on model peak flow (logs)	331	0.258	0.075	0.279	0.550
+ design-capacity control	331	0.036	0.079	0.556	0.550

Notes: System-years (46 systems, 2014–2022); year effects; standard errors clustered by system. Utilization nets design capacity from both sides; the levels slope collapses once design capacity is controlled (last row). The credible-utilization subsample drops systems whose lit-to-potential ratio ever exceeds 1.25.

Appendix D Cable Cutting in History: A New Database and Validation

This appendix presents a validation exercise focused on cable cutting in history. We gather the first systematic database of known cable sabotage events across the first 100 years of the international cable system (1850 to 1945), focusing on cases of military conflict. We stop after WW2 because outright sabotage reportedly ceased during the Cold War, at least in the public record. We then combine the new sabotage data with our traffic-assignment estimates for the historical era. This lets us test whether our modern network approach can explain the targets chosen by belligerent navies in history. The results suggest that it can. A cable's model-derived vulnerability strongly predicts whether it was cut, and belligerents concentrated sabotage on the cables that mattered most for the adversary's connectivity.

The database reveals that cable cutting was particularly pervasive during the two World Wars. Both Great Britain and Germany conducted dozens of cable sabotage operations against each other's cables, across multiple regions such as the Mediterranean, the North Sea, and the Atlantic. This is a novel insight. Britain's cable cutting against Germany during WW1 is a well-known fact recorded in dozens of history books, but the large number of WW2 cutting events was not previously known. To our knowledge, the full extent of cable cutting in WW2 has never before been documented systematically. Against this backdrop, we conduct a case study of cable cutting between Germany and Great Britain in both WW1 and WW2.

Database of cable cutting in history. We compile the most systematic database of cable cutting events assembled to date. For coding, we started with seminal historical works on telegraphy, in particular Winkler (2008), Headrick (1991b), and Kunert (1962), which cover many well-known incidents, especially British and German cable cutting during WW1. In addition, we compile and evaluate primary archival sources, in particular "war diaries" and naval reports, which are near-encyclopedic compendia of thousands of war-related events per year, with comprehensive coverage of both cable cuts and failed attempts. The most important primary sources are (i) a detailed appendix on the Spanish-American War from the US *Report of the Chief of the Bureau of Navigation 1898*; and, for WW2, (ii) the very detailed British *Royal Navy War Diaries* and (iii) Nazi Germany's *Kriegstagebücher der Seekriegsleitung*.

Overall, we identified more than 250 cutting and sabotage incidents targeting international cables since 1850. For several dozen of these, however, we had no primary source or lacked key details, for example on where, when, or which exact cables were targeted. In this version of the database we therefore focus on the three main episodes with census-like coverage: the Spanish-American War of 1898 (19 verified cases), WW1 (67 verified cases), and WW2 (61 verified cases), a total of 147 cable sabotage events. For each event the dataset records the date, the location, the attacker, and the outcome.

Predicting cable cutting. We combine the sabotage data with our historical estimates of each cable's importance for cross-country telecommunication. The main regressor is our model-derived vulnerability proxy, the share of a country's telegraphy traffic estimated to pass through a given cable. We also control for cable length and for whether the cable is a direct link to a wartime adversary.

We exclude the Spanish-American War from the regression, because all international cables around Cuba were targeted and cut within a matter of months, leaving almost no cross-sectional or time variation. The World Wars offer far more variation across space and time, because only about 15 percent of British and German cables were targeted. We focus the analysis on the main "cable war" theaters of WW1 and WW2, namely cable cutting by the UK and Germany in the seas around Europe and the North Atlantic. The sample includes 471 distinct cables controlled by the UK and 412 controlled by Germany.[15] The dependent variable equals one when the adversary cuts the target-country cable in a given year, and zero otherwise. We include only confirmed cuts.

Table D1 estimates a linear probability model by OLS, run separately for German-controlled cables (columns 1 to 4) and British-controlled cables (columns 5 to 8). Cable vulnerability enters positively and significantly in every specification: cables carrying more target-specific traffic, relative to the alternatives, were more likely to be cut. The association is sizable. A one standard deviation increase in traffic on a cable is associated with more than a doubling of the probability that it is cut in a given year.[16] Cable length is likewise positively associated with cutting, while the indicator for a direct link to an adversary is positive but not statistically significant. These results indicate that our model-

[15]"Controlled" here means that at least one of the two landing points lies in the territory of the power in question. When a country was invaded, we reassign its cables to the occupying power. For example, we assume Germany controlled cables landing in the Netherlands, Norway, or Denmark from the date of occupation until each country was liberated.

[16]A one standard deviation increase in traffic is associated with a 1.2 percentage point higher probability of cutting for German cables and a 0.3 percentage point higher probability for British cables (columns 1 and 5, respectively). The unconditional probability of cutting for British cables is only about a third of that for German cables, so the proportional increases are comparable across the two samples.

derived network measures are informative about which cables were cut, with more important cables more likely to be targeted.

Table D1: Predicting cable cutting: Germany and the UK in the World Wars

	German-controlled cables				British-controlled cables			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cable vulnerability	4.493***	3.608***	2.784***	2.031***	2.345***	2.208***	2.992***	2.843***
(own traffic on cable)	(0.351)	(0.276)	(0.518)	(0.549)	(0.886)	(0.839)	(1.096)	(1.042)
Cable length (nm, log)			0.021***	0.021***			0.001***	0.001***
			(0.005)	(0.005)			(0.000)	(0.000)
Direct adversary link			0.070	0.067			0.004	0.005
			(0.044)	(0.042)			(0.007)	(0.007)
Observations	2568	2568	2568	2568	3339	3339	3339	3339
R-squared	0.017	0.049	0.130	0.155	0.004	0.012	0.008	0.015
Conflict FE	Yes		Yes		Yes		Yes	
Year FE		Yes		Yes		Yes		Yes

Dependent variable equals one if the target cable is cut by the adversary in year t. Standard errors in parentheses are clustered by cable. * p ¡ 0.10, ** p ¡ 0.05, *** p ¡ 0.01.