

BAFFI RESEARCH CENTER

Working Paper Master Theses Series

Subsea Cables, Digital Connectivity, and Bilateral Trade

Stefano Ciampa | 2026



**Università
Bocconi**

BAFFI
Centre on Economics,
Finance and Regulation

Subsea Cables, Digital Connectivity, and Bilateral Trade

Stefano Ciampa

May 5, 2026

Abstract

Digital connectivity is increasingly central to international trade. Using bilateral trade flows from the CEPII gravity dataset and data on the global subsea cable network from TeleGeography, I estimate the causal effect of improvements in *bilateral* digital connectivity on bilateral goods trade, exploiting the staggered rollout of new direct submarine cable links. Treatment is defined at the country-pair level as the first year a new direct cable between two countries i and j becomes operational. Because new links can also improve indirect routes for other pairs, I construct a network-based exposure mapping using yearly shortest-path distances to account for interference. I estimate heterogeneity-robust difference-in-differences in a three-way gravity specification. Accounting for spillovers increases the estimated impact; the preferred specification yields an ATT of about 5%. Event-study estimates peak shortly after connection and remain positive for several years. Effects are larger for more distant partners, consistent with improved connectivity reducing bilateral information and coordination frictions and weakening the trade-reducing role of distance.

1 Introduction

Digital connectivity has become an essential premise of the functioning of the international economy. Cross-border transactions increasingly depend on the rapid and reliable transmission

of data to coordinate production, manage supply chains and logistics, match buyers and sellers, deliver services remotely, and verify and settle payments. In parallel, a growing share of the goods and services traded around the globe is either digitally delivered or digitally consumed, making the quality of international connectivity directly relevant for the scale and composition of trade.

The physical backbone of this digital layer is provided by submarine telecommunications cables. Subsea cables are fibre-optic telecommunications cables laid on the ocean floor that connect landing stations across countries and continents, transmitting information as pulses of light through glass fibres. Because these systems offer vastly higher capacity and lower latency than satellite links, as well as substantially lower cost per unit of transmitted capacity, they have become the dominant technology for cross-border data transmission and carry over 99% of international data flows (Mauldin, 2023). By affecting the quality of digital connectivity between countries, these systems can change the effective cost of cross-border trade and information exchange, and ultimately affect bilateral trade patterns even when traditional trade frictions remain unchanged.

Combining data on bilateral trade flows from CEPII with subsea cable data from Telegeography, this paper asks a simple question: *what is the causal effect of an improvement in digital connectivity between two countries on their bilateral trade in goods?* To answer this, I focus on the construction of new direct submarine cable links as a concrete, positive shock to bilateral digital infrastructure between countries. Unlike broader measures of “internet access” at the country level, a direct cable connection is inherently bilateral and therefore maps naturally into a gravity framework. Moreover, this setting naturally lends itself to a difference-in-differences design, but with several important complications. First, treatment adoption is inherently staggered, as country pairs become directly connected in different years, depending on the cable’s Ready-for-Service (RFS) date. Second, the empirical counterpart of the gravity model is inherently non-linear, which makes it natural to work with a non-linear mean specification rather than a linearized model. Taken together, staggered

adoption and non-linearity motivate the use of recent heterogeneity-robust DiD methods developed explicitly for non-linear settings (Wooldridge, 2023; Nagengast and Yotov, 2025). At the same time, identification is challenged by potential violations of the Stable Unit Treatment Value Assumption (SUTVA). This assumption, commonly used in most causal inference studies, posits that each unit’s potential outcomes depend only on its own treatment status and not on the treatment status of other units in the sample. Because the global subsea cable system is a network, countries can benefit from new infrastructure links even when they are not directly connected to it. This means that, when a new link is added, traffic can be rerouted through third countries, and bilateral connectivity may improve for other pairs whose indirect route becomes “shorter” (e.g., in terms of lower latency). In this sense, new cables (the “treatment” of this design) can generate *spillovers*. To account for this treatment interference, I follow an exposure-mapping approach and construct a yearly undirected network where edges correspond to direct cables or terrestrial cables (as proxied by terrestrial contiguity), and countries represent the nodes. A given country pair becomes “exposed” when its shortest-path distance through the network falls below a given threshold. The empirical specification is a three-way ¹ gravity equation estimated by Poisson Pseudo Maximum Likelihood (PPML), and it implements cohort-by-period variation in treatment effects in the spirit of heterogeneity-robust DiD methods.

The results indicate that improved bilateral digital connectivity has economically meaningful effects on trade, and that accounting for spillovers materially increases the estimated impact of a new direct connection. In the preferred specification, the estimated average treatment effect on the treated is about 5.6% percent, and the exposure effect for untreated-but-exposed pairs is also positive and of comparable magnitude. Dynamic estimates show that effects have their peak around the first year after connection, and remain positive for several subsequent years. Additional estimates indicate that the effects are broadly homogeneous across the various cohorts. Lastly, the treatment effect is systematically stronger for more

¹That is, with pair fixed effects and exporter-year and importer-year fixed effects

distant partners, consistent with the idea that improved digital connectivity disproportionately relaxes information and coordination frictions precisely where such frictions are already large. In particular, this last result suggests that improved digital connectivity makes distance less relevant for trade.

2 Related Literature and Contribution

This study relates to several strands of the literature. A first relevant stream analyzes the impact of subsea cable arrivals on internet connectivity and, ultimately, on economic outcomes. [Cariolle \(2021\)](#) studies the arrival of subsea cables in Sub-Saharan Africa and shows that landing countries experience significantly higher internet penetration relative to other countries in the region. Closely related work in the development economics literature exploits subsea cable landings as a source of plausibly exogenous variation in access to fast internet and finds positive effects on firm productivity, firm entry, job creation, and the supply of bank credit ([Hjort and Poulsen, 2019](#); [D’Andrea and Limodio, 2024](#)). Complementary evidence from advanced economies shows that broadband expansion is associated with significant gains in aggregate output and economic performance ([Briglauer et al., 2025](#)). Together, these studies establish that digital infrastructure plays a central role in shaping economic activity, primarily through improvements in local connectivity.

A second closely related stream focuses on cross-border data flows, data protection policies, and international trade ([Casalini and López González, 2019](#); [Ferracane et al., 2023](#); [Chang et al., 2023](#)). Among these contributions, [Ferracane et al. \(2023\)](#) provide direct empirical evidence on the trade effects of data-flow restrictions, showing that countries receiving EU adequacy decisions on data protection experience increases in bilateral digital trade with the EU of roughly 6–14 percent, depending on the definition of digital trade. This study sheds light on the importance of digital frictions as a distinct and increasingly relevant component of trade costs.

This paper is also closely related to the literature on ICT technologies, information frictions, and international trade. Earliest evidence of how ICT technologies affects information frictions is provided by [Steinwender \(2018\)](#), who studies the effects of the laying of the first telegraphic cable between England and the United States on their cotton trade. Her findings suggest that the cable, by significantly reducing communication time, increased trade volumes and allowed cotton traders to better adjust their supply according to the demand conditions in the destination market. Turning to modern technologies, the first study to document the impact of internet on international trade has been [Freund and Weinhold \(2004\)](#). They use the number of web hosts in a country as measure of internet connectivity, and they find that a 10% in the number of web hosts is associated with a 0.2% in the country's export growth. Later studies have further analyzed the effects of connectivity on trade by delving deeper into the underlying mechanisms. [Akerman et al. \(2022\)](#) show that internet adoption can make trade flows more sensitive to distance by lowering information frictions. In their view, better information expands the set of potential trading partners and, ultimately, increases the responsiveness of demand to trade costs. Exploiting firm-level data in France, [Malgouyres et al. \(2021\)](#) study how internet access shapes firms' trading behavior and find that greater internet adoption increases firms' imports, particularly from developing economies. [Fernandes et al. \(2019\)](#), exploiting the rapid diffusion of internet access in China in the pre-Alibaba period, find that higher internet penetration raises firms' probability of exporting by improving information exchange between buyers and sellers. In a similar vein, [Demir et al. \(2024\)](#) show that an upgrade to fast-internet infrastructure affected firms' sourcing patterns within Turkey. In particular, they find that firms reallocate input purchases toward better-connected provinces and diversify their supplier base within those provinces, and they assess that two main channels at play are reduced information and communication frictions. Related work on online platforms shows that digital intermediaries substantially reduce trade costs, promote trade, and facilitate transactions with less familiar and more distant partners ([Carballo et al., 2022](#); [Lendle et al., 2016](#)).

A last stream of literature explicitly studies how changes in *bilateral* digital connectivity affect bilateral trade. [López González and Ferencz \(2018\)](#) estimate that a 10% increase in bilateral digital connectivity raises goods trade by nearly 2% on average, and in a complementary study [López González and Sorescu \(2021\)](#) they find that this effect was about twice as big for trade in parcels. Using a similar definition of bilateral connectivity, [Bellucci et al. \(2025\)](#) find that higher connectivity is associated with lower trade costs, and that these effects are further amplified when effective regulation ensures trust in digital markets and open access to digital infrastructure. More closely related to my specific setting is [Fang et al. \(2024\)](#), who studies how new subsea cables connections affected the exporting behavior of Chinese firms. They find that the firms tended to export 11% more towards directly connected countries and attribute this result to reduced information frictions, more reliable cross-border financial payments, and improved geopolitical relationship. [Imbruno et al. \(2025\)](#) study the extensive-margin effects of submarine cable connections and find that direct links increase the number of exporting firms from high-income countries, while reducing incumbent exporters from poorer economies. They interpret this asymmetry as reflecting lower digital absorptive capacity in low-income countries, so that only the most productive firms can exploit the new opportunities coming from improved connectivity, while less productive incumbents are crowded out.

This paper contributes to these streams of literature in several ways. First, it provides a study on a global scale, in contrast to much of the existing work, which focuses on single countries, regions, or a subset of selected countries, therefore raising concerns about external validity. Second, the analysis is deeply focused on the infrastructure dimension and contributes to the emerging literature that studies subsea cables as bilateral infrastructure that directly increases connectivity between specific country pairs. In fact, some very recent studies adopted this approach ([Imbruno et al., 2025](#); [Fang et al., 2024](#)), but the most influential papers in the literature have treated them mainly as proxies for local internet access. Third, the paper leverages the network structure of the global subsea cable system to characterize countries'

relative positions within the network and to study how new cable connections reshape these positions. This is arguably the main innovation of my study relative to the existing literature, because the network perspective provides a simple and tractable way to account for spillovers to third-country pairs that are not directly connected by the new cable. Empirically, the estimates show that incorporating these spillovers increases the estimated treatment effect. Since prior studies have not modeled this additional channel, the results suggest that existing estimates may be understated, especially in periods of large cable deployments when indirect connectivity improvements are likely to be substantial.

From a methodological standpoint, the paper relates to the literature on robust difference-in-differences designs ([Goodman-Bacon, 2021](#); [de Chaisemartin and D’Haultfoeuille, 2020](#)) and on the estimation of dynamic treatment effects ([Sun and Abraham, 2021](#); [Callaway and Sant’Anna, 2021](#)), as well as their recent extensions to non-linear models and applications to trade data ([Wooldridge, 2023](#); [Nagengast and Yotov, 2025](#)). Finally, the paper builds on recent advances that extend standard difference-in-differences frameworks to settings with SUTVA violations ([Butts, 2023](#); [Fiorini et al., 2024](#)), which are particularly relevant in the context of my study.

3 Background and Data

3.1 Subsea Cables

3.1.1 A brief overview

Subsea (or submarine) telecommunications cables are fibre-optic systems laid on the seabed to carry internet and other data traffic between countries and continents. As discussed in the Section 1, they form the core physical infrastructure of global connectivity and carry the vast majority of international data flows ([Mauldin, 2023](#)). Compared with alternative technologies, like satellite links, subsea cables have become dominant for at least two reasons. First, they

deliver very high transmission capacity at low unit cost. Engineering studies that explicitly compare fibre and satellite show that fibre networks deliver a much lower unit cost and are thus more financially viable and scalable than satellite alternatives ([United Nations, 2017](#)). Second, subsea cables deliver substantially lower latency. The latency gap is largely structural. Geostationary satellites operate at very high altitude, so signals must travel a much longer distance than through terrestrial and subsea fibre routes. As a result, recent engineering measurements consistently report round-trip delays for GEO satellite links in the range of roughly 500–600 ms ([Völk et al., 2024](#); [Deutschmann et al., 2019](#)). By contrast, fibre-based terrestrial and subsea routes follow much shorter paths and therefore sustain substantially lower end-to-end delays. In practice, measurement studies document that median round-trip times on paths that traverse submarine cables can get close to 100 ms for long intercontinental routes, while they stay in the tens of milliseconds for shorter regional links ([Fanou et al., 2020](#); [Livadariu et al., 2024](#)).

A modern subsea cable is a relatively slender object, usually only a few centimeters in diameter for most of its route. At its core are very thin strands of glass fibre, comparable in thickness to a human hair, which guide pulses of light. These fibres are bundled, coated, and then surrounded by successive layers of protective material designed to withstand high pressure and reduce the risk of damage, especially in shallow waters where fishing gear and ship anchors pose a hazard. Cables are manufactured on shore, loaded onto specialised cable-laying vessels, and deployed along carefully surveyed routes on the seabed. In deep ocean they typically lie directly on the seafloor, but the seafloor conditions often play a crucial role in the location decision. Usually, cable routes are chosen to avoid tectonic fault zones, steep continental slopes, rugged seabeds, coral reefs, and areas prone to seismic or volcanic activity, all of which increase the risk of damage ([Pope et al., 2017](#)). For example, the 2006 Hengchun earthquake off Taiwan triggered a massive undersea landslide that broke multiple cables and required 11 repair ships over 49 days to restore connectivity ([Hsu et al., 2008](#)). Notably, human activity also accounts for a large share of cable faults, thus planners seek to

avoid busy fishing grounds and anchorage areas ([TeleGeography, 2026a](#)). In general, whenever the cables are laid nearer to the coast and are more subject to these risks, they are often buried beneath the seabed for additional protection, and then brought ashore only to connect to the landing station.

Indeed, choosing carefully where to place cable routes is quite an important decision for cable-laying companies, as building a new international submarine system involves substantial capital investment. Industry cost estimates indicate that the underwater fibre-optic cable segment (the so-called “wet plant”) alone generally costs tens of thousands of dollars per kilometer, with typical installation costs quoted in the range of USD 30,000–50,000 per km depending on water depth, cable type, and installation conditions ([Straits Research, 2024](#)). Beyond the wet-plant cable itself, projects require extensive planning, marine surveys, regulatory permitting, repeaters spaced along the route, landing stations, and terrestrial backhaul infrastructure, each of which adds further expenditure. Aggregate industry data on new construction show that total annual investment in subsea cable infrastructure has averaged over USD 2 billion globally, with announced projects between 2025 and 2027 alone expected to exceed USD 13 billion in combined value. ([TeleGeography, 2025](#)). Major transoceanic systems have historically involved total project costs in the hundreds of millions of dollars. An exemplary case is that of the West Africa Cable System, which was constructed at a cost of approximately USD 650 million, and the Atlantis-2 cable project cost about USD 370 million.

In terms of durability, subsea telecom cables are engineered with a minimum “design life” of about 25 years for the submerged components (i.e., the wet plant). This design life is defined in terms of acceptable failure rates over that period, rather than a hard cut-off date: some systems have remained physically operational for longer ([TeleGeography, 2026a](#)). In practice, though, some cables may be retired earlier for economic reasons. As technology improves and demand grows, new systems on the same route can offer much higher capacity and lower unit costs, so traffic gradually migrates to them and older cables become commercially obsolete even if they are still technically functional.

3.1.2 Do new subsea cables enhance connectivity?

A key premise behind my empirical design is that the activation of a new *direct* submarine cable between two countries tends to improve their *effective* digital connectivity, especially by expanding capacity and lowering end-to-end latency. This assumption is not innocuous because international connectivity is a complex process shaped by routing policies, congestion, and business relationships. Still, the networking measurement literature provides direct evidence that new cable deployments typically translate into meaningful changes in observed paths and performance, even when alternative routes already exist, thereby leading to increased bilateral connectivity

Measuring the latency effects of a particular cable is difficult in itself, mainly because cable deployments are not observed at the IP layer. Public announcements typically report landing points but not the router interfaces or IP links that actually carry traffic, which makes it hard to connect infrastructures to end-to-end outcomes. Despite these measurement difficulties, several papers have tried to overcome the issues and documented that new direct cables materially affect routing and performance. In this regard, the first detailed study has been provided in [Fanou et al. \(2020\)](#), which analyzes the activation of the South Atlantic Cable System (SACS), the first direct cable connecting Brazil and Angola. The authors find that the new cable reduced round-trip times on relevant paths after activation and conclude that the deployment improved performance for traffic that actually traversed the cable. [Livadariu et al. \(2024\)](#) builds on [Fanou et al. \(2020\)](#) and generalizes this perspective beyond a single case study. They compare end-to-end delays before and after the deployment of 26 subsea cables linking different countries. They find a general improvement in performance across all cable segments, with some showing delay drops by more than 20%. The paper, however, also emphasizes that effects are quite heterogeneous across origin–destination pairs.

Other studies, while not directly measuring the effects of new subsea cables, have documented how subsea cables play a crucial role in ensuring connectivity across countries. [Liu](#)

et al. (2020) show that submarine infrastructure is central to end-user access, documenting substantial dependence of web resource retrieval on submarine paths across a wide set of countries. In addition, Chang et al. (2010) provides evidence from the opposite shock. During their monitoring period, a fault in the SeaMeWe-4 cable that occurred near Italy coincided with severe path congestion, RTT surges, and losses on certain routes because networks carried additional traffic diverted from the faulty cable.

Taken together, these studies support the interpretation that the direct submarine cable is typically associated with meaningful improvements in bilateral digital connectivity, primarily by enabling lower-latency and higher-capacity paths.

3.2 Quality of connectivity and trade in goods: is there a link?

The previous section explained how subsea cables operate and why new cable routes can generate improvements in bilateral digital connectivity. To complete the argument, however, we also need to clarify through which mechanisms increased bilateral connectivity should translate into changes in bilateral trade in goods, which is the main link I am trying to identify in this study. As reviewed in Section 2, the most direct and best-documented channel is that better connectivity reduces information and coordination frictions between trading partners. Several studies have established that faster and more reliable communication facilitates the matching of buyers and sellers, lowers monitoring costs within cross-border relationships, and can ultimately reduce bilateral trade costs (see, e.g., Freund and Weinhold, 2004; Steinwender, 2018; Fernandes et al., 2019; Akerman et al., 2022; Demir et al., 2024; Fang et al., 2024). While this channel provides the primary motivation behind the empirical results obtained in this study, it may not capture the full picture. In fact, when a functioning network already exists (as is the case in my setting), new cables primarily generate *marginal* performance improvements—most notably lower latency—rather than a shift from “no connectivity” to “connectivity”. On the contrary, many of the existing studies rely on big shocks. Steinwender (2018) is the most notable example, as the telegraphic cable she studied was the first ever

laid down in history and allowed information to travel almost in real-time, while before it had to be physically delivered. Similarly, [Akerman et al. \(2022\)](#) analyzes the effects of broadband adoption in Norway, which significantly expanded the set of feasible digital activities compared to the previous dial-up internet. Thus, this subsection proposes additional channels through which improved connectivity could *plausibly* affect bilateral goods trade ².

One way to think about these channels is to go beyond information frictions and communication costs and focus on how connectivity has also *structurally* changed the way goods are traded. In this regard, the diffusion of e-commerce and online platforms represents a major change as it has enabled firms (particularly SMEs) to reach foreign consumers directly, fostering the rise of cross-border parcel trade (i.e. international exchanges shipped as small packages). While existing evidence highlights how these platforms have promoted firm entrance into foreign markets, there are more nuanced margins that can affect firms engaging in international e-commerce. A growing body of evidence suggests that even modest deteriorations in website performance can translate into economically meaningful losses in user activity and conversion rates ³. Using data from an e-commerce platform, [Basalla et al. \(2021\)](#) shows that higher latency worsens key business outcomes along the purchase funnel (with particularly strong effects for mobile users and fast navigators). Similarly, evidence from [Stadnik and Nowak \(2018\)](#) indicates that longer load times for online store pages are associated with lower conversion rates, thus showing that latency affects the probability that users complete a purchase. Also industry studies report evidence in the same direction: small improvements in mobile speed (on the order of tenths of a second) are associated with sizable gains in conversion metrics across major brand websites ([Deloitte Ireland LLP and Fifty-Five, 2020](#)). This evidence also corroborates the results in [López González and Sorescu \(2021\)](#), which finds that increased digital connectivity has an effect twice as strong on parcel trade as

²I cannot test these mechanisms directly using aggregate goods trade. This subsection is presented only as motivation, outlining possible channels through which even small changes in connectivity could affect goods trade, and thus helping to justify what emerges in the results.

³A conversion rate is the share of users (or visits) who complete a target action—e.g., making a purchase—out of all users (or visits) in a given period.

compared to general goods trade. Finally, randomized latency experiments on Google Search show that adding just 100–400ms of delay reduces subsequent user activity (fewer searches per user), with effects that can accumulate over time (Brutlag, 2009). This last study is also consistent with the findings in Fang et al. (2024), which shows that new direct subsea cable connections tend to increase the number of Google searches between connected countries.

Digital technologies have also expanded the set of services that can be traded cross-border. In particular, activities that previously relied on in-person delivery can now be easily delivered remotely. While many of these digitally-deliverable services are fairly tolerant to latency (e.g., asynchronous messaging, file transfer) and are, thus, independent from the quality of connectivity, several relevant categories are *latency sensitive* because they rely on synchronous interaction, tight feedback loops, or stable real-time rendering. Classic Quality of service (QoS) guidance distinguishes interactive from non-interactive applications by highlighting the level of delay that is tolerable to ensure a smooth delivery (International Telecommunication Union, 2001, 2003). Latency-sensitive services are typically those that require smooth, real-time interaction across borders (such as live video calls, remote support or troubleshooting, and some advanced digital applications used for machine automation) where delays make communication and coordination noticeably harder (International Telecommunication Union, 2001, 2003; ETSI, 2024, 2022; 3GPP, 2023). Although these mechanisms explicitly concern services, they can also (indirectly) affect trade in goods. In fact, in recent years, firms have been increasingly selling integrated solutions in which the good is complemented by (often remotely deliverable) services provided before and after the sale, a trend that is often called "servicification" of merchandise trade (UNCTAD, 2025). In fact, some goods, such as complex machinery, electrical, and ICT equipment, require services such as maintenance or automation that are often latency-sensitive (particularly automation) (Cadestin and Miroudot, 2020; Miroudot and Cadestin, 2017; Miroudot, 2017).

Overall, existing evidence suggests that even small reductions in bilateral latency, such as those induced by a more direct subsea route, can strengthen bilateral trade through two

complementary channels: they can raise the competitiveness of platform-mediated e-commerce by improving user experience and platform performance metrics, thereby increasing digitally ordered parcel trade, and they can improve the quality of latency-sensitive service provision bundled with goods, which enhances the attractiveness of integrated solutions and ultimately increases demand for the underlying products.

3.3 Data and Descriptives

3.3.1 Datasets used

Data on bilateral trade flows are drawn from the CEPII Gravity database ([Conte et al., 2022](#)), which provides bilateral trade data and standard gravity covariates for all country pairs up to 2020. I then merge these data with information on the global subsea cable network from TeleGeography. The cable data report each system’s Ready For Service (RFS) date, length, landings, and covers the universe of active (or under-construction/planned) cable systems as of 2022.

A key limitation is that the observed network is not a complete historical census, as the set of cables recorded at a given RFS date reflects systems that are still active in 2022, while older systems with the same RFS date may have been decommissioned and thus do not appear. This implies that, although some RFS dates in the data go back to the early 1990s, the resulting picture is inevitably less reliable for the earliest years, when decommissioning is more likely to have occurred. A second limitation is that the dataset does not provide information on terrestrial fiber networks, which constrains what can be said about infrastructure-driven changes in connectivity for landlocked countries, and also complicates the definition of spillovers. In fact, country pairs may benefit indirectly from new subsea capacity through terrestrial backhaul and cross-border terrestrial links, but these channels cannot be measured *directly* with the available data. A further limitation of the cable dataset is that it includes detailed attributes such as Ready For Service (RFS) dates and cable lengths, but provides

capacity-related information only for a few cable systems. The more granular metrics on lit and potential capacity and fibre-pair counts are tracked by TeleGeography’s subscription research services and are not available in the publicly accessible dataset ([TeleGeography, 2026a](#)).

3.3.2 The evolution of the subsea cable network.

Figure 1 reports the distribution of newly built cable routes by country pair, for each different RFS year ⁴, and shows how cable construction is highly clustered around specific years. The first prominent spike occurs in the early 2000s, coinciding with the late dot-com era, when cable deployment accelerated rapidly, especially along major long-haul routes. Contemporary industry accounts highlight that this surge led to substantial oversupply and a subsequent “shakeout,” marked by bankruptcies and large write-offs among operators ([Winseck, 2017](#); [Warf, 2011](#)). Thus, it is very likely that, in the immediate years after their deployment, the cables were largely underutilized and did not lead to significant improvements in digital connectivity between countries. A second investment wave emerges from around 2010 onwards. Unlike the dot-com bubble, this phase is most likely driven by sustained growth in global data demand and the increasing involvement of large content and cloud providers, which began investing directly in cable systems rather than relying solely on wholesale capacity ([Caoui and Steck, 2025](#)).

Figure 2 illustrates instead the geographic expansion of the global subsea cable network between 2000 and 2020. In 2000, the network is sparse and highly concentrated, with a main long-haul corridor already linking North America and Western Europe, while the Middle East and parts of East Asia appear connected through a limited set of routes. By contrast, Sub-Saharan Africa and most of South America were still largely peripheral, with very few connections to the main backbone.

⁴Basically, it accounts to the number of country pairs that get directly connected within the newly built cable systems.

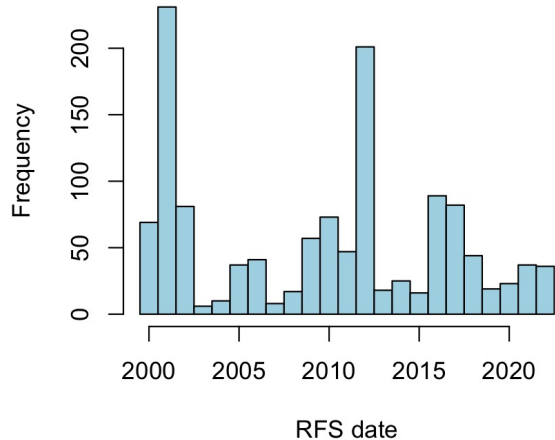


Figure 1: Newly added direct cable routes (by country pair)

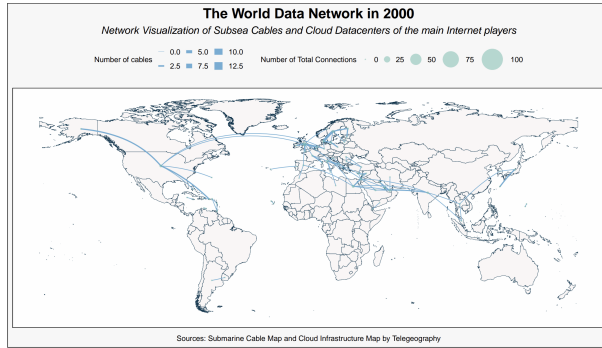
By 2005, connectivity becomes visibly denser. The transatlantic corridor thickens as additional cable routes are built, and Asia appears more capillarially connected, reflecting a rapid expansion of both intra-regional and inter-regional links. At the same time, previously lagging regions begin to catch up. West Africa starts to connect more systematically to Europe, and South America strengthens links primarily to North America. One notable development is the emergence of Egypt as a key global hub. Much like in maritime trade, the Suez Canal corridor represents a major bottleneck for cable routes connecting Asia and Europe, which has allowed Egypt to occupy a central brokerage position in the network. To give a sense of magnitude, in 2005 Egypt was directly connected to 30 countries, compared with 24 for the United Kingdom.

From 2010 onward, the maps mostly reflect progressive densification and increasing redundancy in these lagging regions. Africa gains additional connections along both western and eastern routes, and Latin America becomes more connected not only through North American gateways but also through a broader set of international routes. By 2015 and especially 2020, the network looks global and significantly more interconnected, with multiple

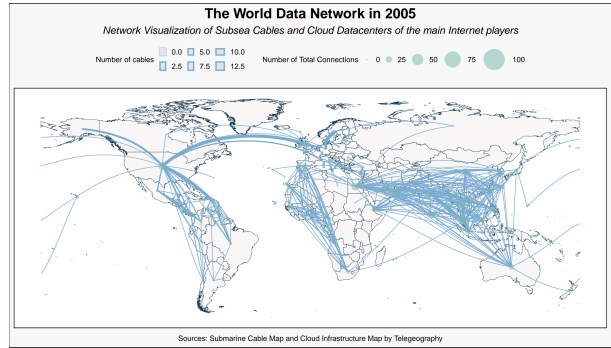
alternative paths across major corridors and the emergence of large hubs around key landing points ([TeleGeography, 2026b](#)). As of 2020, Egypt is the most connected country, with direct links to about 93 countries, followed by Malaysia and India, with roughly 80 and 73 direct connections, respectively. Western countries also rank highly. Most notably, the United States is directly connected to about 68 countries, and France to about 47.

3.3.3 Sample Selection and Descriptives

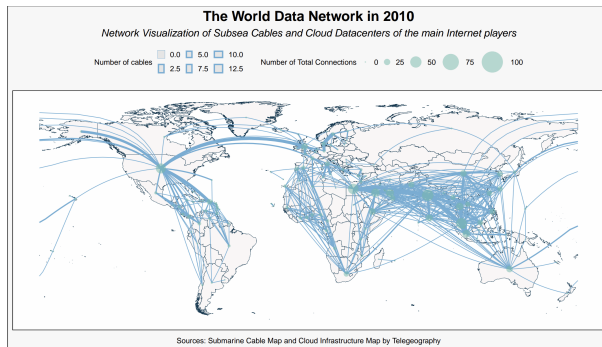
The sample period chosen spans from 2000 to 2020, where 2020 is the last year available in the CEPII gravity dataset. However, among treated country-pairs, I restrict attention to those that received treatment in 2005 or later, up to the 2019 cohort. This restriction is motivated by two considerations. First, because the cable data exclude dismantled systems, the observed network is subject to survivorship bias since, for any given RFS date, only cables that are still active today are observed. This implies that for very early cohorts, the estimated treatment effect may be severely biased. Focusing on more recent RFS dates, therefore, reduces the likelihood of missing dismantled cables and gives more credibility to the identification of cohort-level treatment effects. In fact, under this restriction, the only omitted systems are those that were constructed and dismantled between 2005 and 2020, which, given the typical lifespan of subsea cables ([TeleGeography, 2026a](#)), should represent only a limited share of all relevant connections. Second, the early 2000s were characterized by a period of substantial oversupply following the dot-com boom in subsea cable investment, as compared to the real demand for data at the time. Thus, given that the effect I am interested in is precisely that of an increase in bilateral connectivity, including this period in the main analysis would likely attenuate the estimated effects, particularly in the years immediately following connection, as much of the newly installed capacity was initially underutilized. In the appendix, I report estimates that include earlier cable connections and show that the results are qualitatively similar, although the point estimates are somewhat smaller. Finally, I restrict the analysis to country pairs with non-missing trade data throughout the 2000–2020 period, so that the



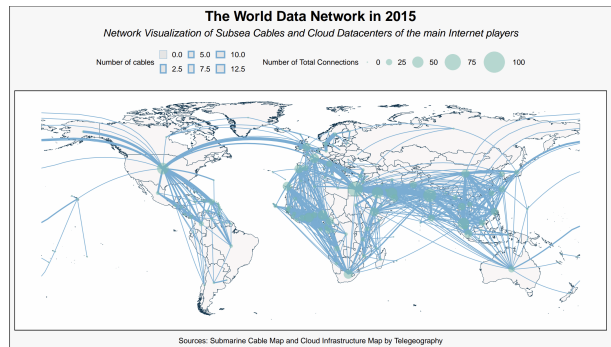
(a) 2000



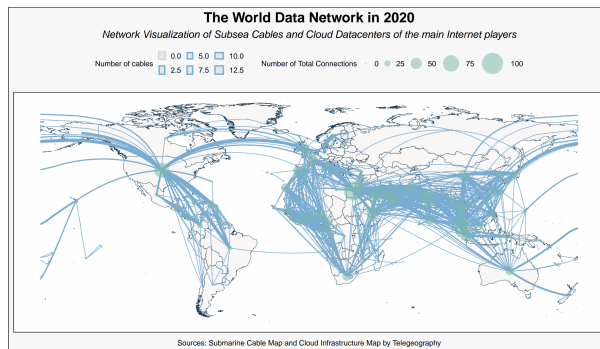
(b) 2005



(c) 2010



(d) 2015



(e) 2020

Figure 2: Network evolution, 2000–2020.

resulting panel is balanced. While this restriction is not strictly required for identification, it improves the precision of the estimated pair fixed effects, thereby reducing noise in the estimated treatment effects.

Looking at the sample’s descriptive statistics, table 1 shows some compositional differences between treated and control pairs. Control pairs are, on average, farther apart in terms of physical distance. This likely reflects that, despite the construction of some very long transoceanic systems, most subsea systems connect relatively nearby landing points and primarily strengthen intra-regional connectivity. Treated pairs also display systematically higher values for other proximity measures (contiguity, colonial ties, common language and legal origin, and religious proximity), which is unsurprising given that geographic closeness tends to correlate with these forms of historical and institutional similarity. By contrast, the two groups look very similar in terms of exporter and importer GDP growth and population growth, suggesting no major differences in aggregate macro growth trends. Finally, trade growth is somewhat higher among control pairs. This is consistent with recent findings in the literature, which document that long-distance trade has increased over the past decades (Ganapati and Wong, 2023; UNCTAD, 2024; World Economic Forum, 2021)

Table 1: Descriptive statistics

Variable	Control: Mean (SD)	Treated: Mean (SD)	Diff.
Distance (km)	7311 (4392)	3354 (2694)	-3957
Contiguity indicator	0.025 (0.155)	0.105 (0.307)	0.080
Colonial link (ever)	0.166 (0.372)	0.301 (0.460)	0.135
Common official language	0.153 (0.360)	0.288 (0.454)	0.135
Common legal origin (pre-1991)	0.310 (0.463)	0.559 (0.498)	0.249
Religious proximity (0–1)	0.166 (0.244)	0.231 (0.253)	0.065
Exporter GDP growth	0.063 (0.028)	0.063 (0.029)	-0.001
Importer GDP growth	0.064 (0.029)	0.063 (0.029)	-0.001
Exporter population growth	0.012 (0.013)	0.015 (0.013)	0.003
Importer population growth	0.013 (0.013)	0.016 (0.014)	0.003
Trade growth	0.082 (0.109)	0.058 (0.095)	-0.024

Notes: Growth rates are computed as the geometric average over the sample period. $N_{control} = 13483$, $N_{treated} = 229$. SD in parentheses.

4 Empirical Strategy

4.1 Treatment and Exposure: Definitions and Assumptions

As discussed in the Introduction, the treatment I want to study is an increase in the quality of the bilateral digital infrastructure, as proxied by the laying of a direct subsea cable connection between the two. For each pair (i, j) , let RFS_{ij} denote the year in which a subsea cable between i and j becomes ready for service (RFS). Treatment is then given by the indicator

$$d_{ijt} = \mathbf{1}\{t \geq RFS_{ij}\},$$

so that $d_{ijt} = 0$ in all years before the cable is operational and $d_{ijt} = 1$ in all subsequent years. Country pairs that are never directly connected by a cable in the sample period are never treated and thus have $d_{ijt} = 0$ for all t . Given that each pair has a different RFS year, this setting naturally gives rise to a staggered-adoption design as different country pairs are treated in different years. For simplicity, I will use g instead of RFS to index the treatment cohort in the next section⁵.

A key challenge for identification in this setting is the presence of spillovers across country pairs, as new subsea cables can generate benefits beyond directly connected country pairs. Under the Stable Unit Treatment Value (SUTVA) assumption, the potential outcome of a given pair (i, j) would depend only on its own treatment status, that is, on whether there is a direct subsea cable between i and j . This assumption is not realistic here. Even if (i, j) is not directly connected by a new cable, both countries may gain from improved indirect routes. For example, if country i is directly connected to k and country j also is directly connected to k , improvements in digital connectivity along the paths $i-k$ and $j-k$ may still generate benefits between i and j . As a result, the trade outcome for pair (i, j) may also depend on the broader pattern of cable connections involving their neighbors. With spillovers, outcomes

⁵This is to maintain my notation consistent with existing literature on staggered difference in differences

may depend, in principle, on the entire treatment vector $D_t = (d_{klt})_{k,l}$, so that each unit has 2^N potential outcomes, where N is the number of units.

To make this problem tractable, I follow the exposure-mapping approach and summarize the full treatment vector into a lower-dimensional exposure variable. Formally, an exposure mapping is a function

$$h : \mathcal{Z} \rightarrow \mathcal{X},$$

where \mathcal{Z} is the set of all possible treatment vectors and \mathcal{X} is a researcher-defined set of exposure conditions. The key idea is that, even though there are many potential outcomes, we can define subgroups of them that are in some way “similar”, which we can thus synthesize in a few conditions. In this paper, I specify a simple exposure mapping that classifies each country pair as either “exposed” or “not exposed” to spillovers from subsea cables.

To understand how, in my setting, two countries may be exposed and benefit from the construction of another direct subsea cable connection, it is important to have at least a vague idea of how international routing works. When two countries are not directly connected, and thus have to rely on different cable systems to reach one another, it usually entails more “handoffs” across networks between the various operators, and therefore frictions. In particular, Internet routes can be substantially longer than a direct physical path due to intra- and inter-domain policy choices and peering constraints (the so-called “path inflation”) ([Spring et al., 2003](#)). Studies have found that longer transit paths correlate with higher latency (as measured by the RTT) and exhibit higher propagation and queueing delays ([Huffaker et al., 2002](#); [Ahmed et al., 2017](#)). Also, even when end-to-end physical capacity is ample, bottlenecks can arise at interconnection links between Internet Service Providers, where persistent congestion episodes have been documented ([Clark et al., 2014](#)). Thus, it is reasonable to define an exposure mapping in which country pairs that can reach each other through fewer cable-system traversals are expected, on average, to face fewer frictions and can thus benefit more from improved connectivity in any segment of the path.

To operationalize this idea, I adopt a network-based approach. I construct, for each year, an unweighted and undirected graph in which countries are nodes and edges indicate either terrestrial contiguity or the presence of a subsea cable between two countries. Terrestrial contiguity is used to proxy for terrestrial backbone networks, which are not observed in the dataset. Given this yearly network, I compute the shortest-path distance between every pair of countries using Dijkstra’s algorithm. As more cables are built over time, the network becomes denser, new “shortcuts” appear, and the shortest distance between many country pairs decreases. A country pair enters the “exposed” status ⁶ in the first year in which its shortest-path distance falls to four hops⁷ or less, which is defined as EXP_{ij} . Since this cutoff is admittedly arbitrary, I re-estimate the model using alternative distance thresholds and show that the main results are robust to these choices. In addition, I also assume that countries that are already treated are not exposed to spillovers and forfeit any previously received spillovers ⁸. Exposure is thus given by the indicator:

$$e_{ijt} = \mathbf{1}\{EXP_{ij} \leq t < RFS_{ij}\}$$

4.2 Potential Outcomes Framework

4.2.1 Definitions

Given the definitions of treatment and exposure above, it is convenient to work with treatment and exposure histories. For each country pair (i, j) , let

$$\underline{d} = [d_1, \dots, d_T] \quad ; \quad \underline{d}^t = [d_1, \dots, d_t],$$

⁶Indeed, since this methodology relies on a proxy rather than direct observation of terrestrial links, it would be more appropriate to refer to “potential exposure.” However, for simplicity, I will continue to use the term “exposure”.

⁷In network theory, the number of hops between two nodes is the number of edges traversed along a path connecting them.

⁸See Assumption 3

where \underline{d} denotes the full treatment history over the sample period, \underline{d}^t the history up to time t , and $d_t \in \{0, 1\}$ is the binary treatment indicator at time t . Since each cable in the dataset remains active until the end of the sample, I assume irreversibility of treatment:

$$d_{t+1} \geq d_t \quad \forall t.$$

Spillovers from other country pairs' treatment statuses are captured through exposure. The associated exposure histories are:

$$\underline{e} = [e_1, \dots, e_T] \quad ; \quad \underline{e}^t = [e_1, \dots, e_t],$$

Where $e_t \in \{0, 1\}$ is the binary exposure indicator at time t .

4.2.2 Identifying Assumptions and the quantity of interest

The identifying assumptions I rely on are standard in the difference-in-differences literature. However, given the specific framework of my analysis, it is useful to state them explicitly, one by one, to facilitate interpretation and make their content easier to visualize.

The first identifying assumption is no anticipation:

Assumption 1 (No anticipation). *Let $Y_{ijt,g}(\underline{d}, \underline{e})$ denote the bilateral trade flow between exporter i and importer j at time t under treatment history \underline{d} and exposure history \underline{e} . Then*

$$Y_{ijt,g}(\underline{d}, \underline{e}) = Y_{ijt,g}(\underline{d}^t, \underline{e}^t).$$

This assumption implies that potential outcomes at time t depend only on treatment and exposure realized up to t , not on future realizations. In particular,

$$Y_{ijt,g}(\underline{d}^t, \underline{e}^t) = Y_{ijt,g}(\underline{0}^t, \underline{0}^t) \quad \forall t < \min\{EXP, RFS\},$$

so that, prior to the first exposure or treatment date, the relevant potential outcome coincides with the Never-Exposed and Never-Treated (NENT) potential outcome.

Then, I assume conditional parallel trends:

Assumption 2 (Conditional Parallel Trends). *Let $Y_{ijt,g}(\underline{d}^t, \underline{e}^t)$ denote the bilateral trade flow between exporter i and importer j at time t under treatment history \underline{d}^t and exposure history \underline{e}^t . Let $Y_{ijt,g}(\underline{0}^t, \underline{0}^t)$ denote the NENT potential outcome. Then there exists a set of fixed effects $\{\alpha_{it}, \gamma_{jt}, \eta_{ij}\}$ and a vector of time-varying covariates \underline{x}_{ijt} such that*

$$\mathbb{E}\left[Y_{ijt,g}(\underline{0}^t, \underline{0}^t) \mid \eta_{ij}, \alpha_{it}, \gamma_{jt}, \underline{x}_{ijt}\right] = \exp\left(\eta_{ij} + \alpha_{it} + \gamma_{jt} + \underline{x}'_{ijt}\underline{\beta}\right).$$

Relative to a standard DiD design, the conditional parallel trends (CPT) assumption used here requires parallel evolution of trade flows only after conditioning on exporter-specific, importer-specific time shocks and differential trends related to a set of observable time-varying covariates. In particular, allowing for importer-specific or exporter-specific trends implies that the effective treated-control comparison is happening in a more nuanced way with respect to the usual DiD designs. To understand where identification is coming from, consider countries i and j and two comparison partners k and n , the ratio of normalized NENT trade flows is ⁹:

$$\frac{\tilde{Y}_{ijt}^0 / \tilde{Y}_{int}^0}{\tilde{Y}_{kjt}^0 / \tilde{Y}_{knt}^0} = \exp\left(\tilde{\mu} + \tilde{\underline{x}}'_t \underline{\beta}\right)$$

where $\tilde{\mu} = (\mu_{ij} - \mu_{kj}) - (\mu_{in} - \mu_{kn})$ is time constant and $\tilde{\underline{x}}'_t = (\underline{x}_{ijt} - \underline{x}_{kjt}) - (\underline{x}_{int} - \underline{x}_{knt})$ is time-varying but known. Thus, to assess the effect of a hypothetical cable between i and j , the model would effectively compare the export growth from i to j relative to the export growth from i to n , with the export growth from k to j relative to k to n , accounting for the effects of the covariates on their respective trends. This ratio-based intuition mirrors the logic behind the methods used by [Head and Ries \(1999\)](#) and [Romalis \(2007\)](#) to estimate the

⁹For simplicity, I have set $\tilde{Y}_{ijt}^0 = \mathbb{E}[Y_{ijt,g}(\underline{0}, \underline{0}) \mid \eta_{ij}, \alpha_{it}, \gamma_{jt}, \underline{x}_{ijt}]$

impact of trade agreements without explicitly solving for multilateral resistance terms.

Another assumption, specific to this setting with SUTVA violation, is the No Spillovers on the Treated (NST) assumption, which is formally expressed as:

Assumption 3 (No spillovers on the treated). *Let $Y_{ijt,g}(\underline{d}^t, \underline{e}^t)$ denote the bilateral trade flow between exporter i and importer j at time t under treatment history \underline{d}^t and exposure history \underline{e}^t . Then, for all $t \geq RFS$,*

$$Y_{ijt,g}(\underline{d}^t, \underline{e}^t) = Y_{ijt,g}(\underline{d}^t, \underline{0}^t).$$

Imposing this assumption yields practical advantages from an econometric standpoint. Relaxing it in full generality would require identification of the treatment effect under different sets of exposure histories, leading to the estimation of a potentially large number of interaction terms between the treatment indicator and indicators for different patterns. In settings with many possible exposure profiles, the problem quickly becomes high-dimensional and can lead to a substantial reduction in precision.

However, it is crucial to highlight that this assumption has two key implications for the interpretation of the estimates. The most obvious and intuitive one is that it rules out any *post-treatment spillovers* affecting treated pairs (i.e., the “ex post” implication). In other words, once a country pair becomes treated, the model assumes that subsequent changes in the surrounding network do not generate additional effects on that treated pair beyond the treatment itself. This is fairly reasonable in this context since, once a direct link is in place, the main reduction in routing frictions has already occurred, and further improvements in indirect paths are likely to translate into only limited gains. Second, the assumption also imposes homogeneous treatment effects with respect to pre-treatment exposure (i.e., the “ex ante” implication). That is, the estimated treatment effect is assumed to be the same irrespective of whether a pair had already experienced improvements in indirect routing before receiving a direct cable. This is a more stringent requirement. If a pair already enjoyed great improvement in indirect connectivity, the marginal benefit of adding a direct link may be

smaller. By contrast, for pairs facing weak or fragmented indirect connectivity ex ante, the introduction of a direct cable could represent a much larger discrete improvement, implying a larger incremental effect.

Testing the ex-post implication is not possible, given the way the exposure mapping is constructed. Once a country pair is directly connected, its network distance is mechanically equal to one, and further changes in the network cannot reduce this distance. Thus, additional indirect routes cannot generate measurable spillovers for already-treated pairs. By contrast, the ex ante implication is testable. To assess this, I augment the baseline specification with an interaction between the treatment indicator and an indicator for ex ante exposure, so that the main treatment effect is identified from pairs that were never exposed before receiving a direct cable. The results show that the implied average treatment effect on pairs that were never exposed before is higher (although not statistically significantly different from) the baseline estimate. This, however, signals that this additional margin of heterogeneity may be relevant to assess the benefits of a new direct cable connection.

Given the identifying assumptions, the quantity of interest is the cohort-time average treatment effect on the treated (in percentage change) under no interference, which is formally defined as:

$$ATT_0^{\%}(g, t) = \ln \left(\frac{\mathbb{E}[Y_{ijt,g}(\underline{d}^t, \underline{Q}^t) \mid \eta_{ij}, \alpha_{it}, \gamma_{jt}, \underline{x}_{ijt}]}{\mathbb{E}[Y_{ijt,g}(\underline{Q}^t, \underline{Q}^t) \mid \eta_{ij}, \alpha_{it}, \gamma_{jt}, \underline{x}_{ijt}]} \right)$$

Intuitively, this quantity represents the effect of improved bilateral connectivity while no other country in the network. To better visualize it, it is as if the two countries were isolated islands in the middle of the ocean. Once they get connected, their bilateral connectivity improve but no other third country benefits from the new link. However, estimating this quantity even for the latest cohorts requires that a non-trivial share of NENT units remain in that status until the end of the sample. This may sound like an impossible requirement to satisfy, as most coastal countries are eventually connected to the network by the end of the sample. However, the key effect I want to analyze is that of a *sudden* improvement in

bilateral connectivity and, even though most countries are connected, some of them lie at a considerable network distance from one another and thus their data has to travel several different cable systems to reach the other partner. In my data, given the definitions and assumptions above, approximately 30% of all observations are classified as NENT.

4.3 Empirical specification and estimation

For the estimation, I adopt the approach in [Wooldridge \(2023\)](#), which extends robust diff-in-diff methods with staggered treatment to non-linear models. In line with the gravity literature, the model is estimated by Poisson pseudo-maximum likelihood (PPML) with standard errors clustered at the pair-level. PPML is the workhorse estimator for structural gravity equations because it naturally handles heteroskedasticity and allows to retain zero trade flows. In my setting, I augment the model to account explicitly for spillover (exposure) effects, and I allow the treatment effect to be heterogeneous with respect to log distance and the (five-year) pre-treatment average log GDP of both the importer and the exporter. I use pre-treatment (rather than contemporaneous) GDP to avoid bad controls; that is, covariates that are time-varying and that could themselves respond to treatment. In fact, the arrival of a subsea cable can have positive effects on a country's economy and, trivially, if the cable truly has a causal effect on bilateral trade, part of that effect will naturally be reflected in contemporaneous GDP. Using pre-treatment averages, therefore, helps ensure that the variables used to capture heterogeneity are predetermined with respect to the treatment. Lastly, I control for whether the two countries are engaged in a regional trade agreement.

The final empirical specification takes the following form:

$$y_{ijt,g} = \exp \left(\begin{array}{l} \boxed{\mu_{ij} + \alpha_{it} + \gamma_{jt} + \beta RTA_{ijt}} \quad (1) \text{ Baseline outcome} \\ + \boxed{\sum_{g'=2005}^{2019} \sum_{t'=2005}^{2019} \mathbb{I}(g' = g) \mathbb{I}(t' = t) (\delta_{g,t} + \tilde{z}'_{ijt} \eta_{g,t})} \quad (2) \text{ Treatment effect} \\ + \boxed{\sum_{t'=2000}^{2019} \mathbb{I}(EXP \leq t < g) \mathbb{I}(t' = t) \zeta_t} \quad (3) \text{ Exposure Effect} \end{array} \right)$$

where, for notational convenience,

$$\tilde{z}'_{ijt} = [\ln(Dist_{ij}), \ln(GDP_i), \ln(GDP_j)]$$

The term in (1) captures the NENT outcome, that is, the outcome for a pair that is never treated and never exposed. The term in (2) is the generalized treatment effect, varying flexibly by cohort g and year t and interacting with the chosen treatment heterogeneity variables, and the term in (3) captures the effect of being exposed while being untreated. Each of the interacted covariates is demeaned at the cohort-time level, so that the main treatment effects $\delta_{g,t}$ can be interpreted as the average effects on the treated. In the results section, I also show aggregate treatment effects as well as the dynamic event-study estimates. It should be noted, however, that all the estimated coefficients should be interpreted as partial-equilibrium effects. If a new subsea cable reduces trade frictions, the implied adjustment would also affect multilateral resistance terms. A general-equilibrium assessment would therefore require imposing the full set of equilibrium conditions in the estimation (Yotov et al., 2016), which is something that goes beyond the scope of this study.

I also relate this robust specification to the simpler “two-way” diff-in-diff approach, in which the treatment effect is summarized by a single treatment dummy rather than a full set of cohort-by-period interactions. In fact, unlike the robust estimator, the simple two-way estimator yields a coefficient that is a weighted average of all possible 2×2 diff-in-diff

comparisons, some of which are “forbidden” because they use already-treated observations as controls. Intuitively, the size of the resulting bias depends on how frequent such comparisons are (de Chaisemartin and D’Haultfoeuille, 2020). In my application, treated observations account for only about 6% of the sample, which mechanically limits the incidence of forbidden comparisons. Accordingly, I expect the simple two-way estimator to be close to the robust cohort-by-period estimate, and the empirical results confirm that the two sets of estimates are indeed very similar.

In addition to the main effect, the interaction term with distance is of particular interest. On the one hand, it can be read as the “additional” effect of a new cable connection as distance varies. In this view, a positive interaction implies that the gain from improved bilateral connection is larger for more distant partners. On the other hand, it can be interpreted as the effect of the cable connection on the distance elasticity of trade, even though the level effect of distance is absorbed by the pair fixed effects. This margin of heterogeneity can shed useful insights about the role of digital connectivity in reducing information and communication frictions. In fact, more distant pairs usually incur higher frictions, and thus we may expect these countries to benefit more. A similar logic applies to the interactions with GDP. Existing studies have highlighted how developing economies strongly benefit from increased internet connectivity (Hjort and Poulsen, 2019; D’Andrea and Limodio, 2024), thus we may expect the effect of the cable to be stronger for them.

4.3.1 The issue of countries connected by multiple cables

It is not uncommon for the same country pair to get connected several times. This may be due to several reasons. For instance, it may be that the older cable has become partly obsolete and gets complemented with a newer, more up-to-date cable, or simply it may be that as data demand grows, the need for additional capacity calls for a new cable to be built. In any case, these additional cables improve bilateral connectivity, and it may thus be appropriate to include them in the analysis. In fact, pairs that are connected more than once

represent about half of the treated sample and thus account for a substantial share of treated observations. To address this, I treat each cable event separately for a given country pair and then reweight the estimated effects by the total number of cables that connect that pair over the sample. In practice, this implies that pairs with multiple connections contribute to the ATT through the average effect of all cables that have ever connected them. Although this procedure is somewhat unconventional, it has several advantages. First, given the already limited number of directly connected pairs, it allows me to retain a large treated sample, instead of eliminating multiply connected observations. Second, choosing a single cable for pairs with multiple connections (e.g., the first or last) would make the event–study coefficients difficult to interpret, since pre- and post-treatment periods would still be affected by other cable connections that are intentionally excluded from the sample. This last problem does not disappear under the reweighting approach but it is attenuated, as for pairs with multiple connections, periods before and after a given cable get averaged, so the influence of other cables on the estimated dynamic pattern is reduced.

An alternative approach would be to interpret the number of cables as a treatment intensity and estimate a model with multiple treatment levels. However, the methodological literature does not yet provide a robust framework for staggered diff-in-diff with units that transition across several treatment intensities over time (Callaway et al., 2024). In my context, any pair connected at least twice would span multiple treatment levels (i.e., no cable, one cable, two cables), making a standard intensity framework difficult to implement. Moreover, such an approach would be most useful if the goal were to study the marginal effect of additional cables and, for example, to test for diminishing returns to new connections. While this is an interesting question, both data limitations and the nature of the technology make it challenging in practice. First, my dataset does not contain information on individual cable capacities, so I cannot make clear assessments on the quality of the new cables. Second, even with detailed capacity data, the subsea cable industry has undergone substantial technological change, with design capacities increasing by an order of magnitude or more over the past two

decades, and it is unclear how to map such changes in engineering characteristics into the evolving computational and data needs of digital technologies that may affect trade. Making precise statements about marginal returns would therefore require not only richer data but also a much deeper technical understanding of the underlying systems.

Instead, my analysis relies only on two relatively robust facts: (i) whenever a direct connection is built, it typically improves digital connectivity between the two connected countries or regions, even when indirect alternative routes already exist; and (ii) additional direct cables—beyond the first—are typically deployed to augment existing connectivity (e.g., by increasing capacity, redundancy, or reliability), so it is reasonable to assume that they improve connection quality relative to the previously installed infrastructure. . The first fact justifies using subsea cable construction as the treatment variable, and the second justifies allowing subsequent cable constructions for the same country pair to affect the quality of their connectivity, while remaining agnostic about whether the marginal gain in connectivity increases or decreases with the number of cables.

Nonetheless, to show that this procedure does not materially affect the results, I also conduct a robustness check. In the appendix, I replicate the main analysis by restricting attention, for each country pair, to the first cable laid after 2005, and I show that the findings are broadly unchanged and consistent with the baseline specification.

5 Results

5.1 Diff-in-diff estimates

Table 2 reports the diff-in-diff of both the robust and ordinary TWFE estimates. When spillovers are ignored, the estimated average treatment effect of a direct cable connection is positive but not statistically significant. Once spillovers are explicitly accounted for, the estimate of the main effect increases, becomes statistically significant at the 5% level, and the

spillover coefficient is also positive and of similar magnitude. This indicates that country pairs that do not receive a direct cable but become exposed through the network also experience higher bilateral trade, and thus not accounting for such spillovers can lead to underestimating the true treatment effect. Even though the exposure coefficient cannot be strictly interpreted as causal ¹⁰, this result adds robustness to the exposure framework, as it shows that both the direct treatment effect and the spillover effect are positive, which is logically consistent with the idea that improved connectivity (whether direct or indirect) raises bilateral trade.

The interaction between treatment and distance is positive and statistically significant, both in specifications that ignore spillovers and in those that account for them. This implies that the trade gains from a new cable are larger for geographically more distant pairs. Equivalently, since distance has a negative impact on trade in the baseline gravity structure ¹¹, these results suggest that cable connections reduce the magnitude of the distance elasticity, making geographic distance less detrimental for bilateral trade among connected countries. Since distance proxies not only for transport costs but also for information and coordination frictions, the result is also consistent with improved digital connectivity relaxing these frictions precisely where they are ex ante more severe. Overall, these results are consistent with the large literature documenting how internet diffusion caused the “death of distance” (Carballo et al., 2022; Lendle et al., 2016).

On the other hand, the interactions with GDP are negative, although only weakly significant for destination GDP. Nonetheless, the sign of the coefficient is broadly in line with evidence that connectivity improvements yield larger gains in poorer economies (D’Andrea and Limodio, 2024; Hjort and Poulsen, 2019), suggesting diminishing marginal benefits of connectivity in richer and more advanced economies.

Comparing these estimates with the ordinary TWFE specification, the estimated ATT is very similar across the two approaches, although the robust estimator yields a slightly

¹⁰To prove causality in this case, we would need to ensure conditional parallel trends also for exposed vs unexposed units.

¹¹even though distance is not explicitly accounted for in the estimation as it is absorbed by pair-fixed effects

larger effect, both for the direct treatment and for the exposure coefficient. This pattern is consistent with the idea that, in my setting, the bias of the ordinary TWFE estimator is limited because of the small share of treated observations in the sample.

Table 2: Aggregate treatment effects

	Dependent variable: $Trade_{ij}$			
	TWFE	TWFE	Robust TWFE	Robust TWFE
$Connected_{ij}$	0.0170 (0.0166)	0.0393** (0.0189)	0.0267 (0.0170)	0.0556*** (0.0195)
$Connected_{ij} \times \log(Dist_{ij})$	0.0267** (0.0117)	0.0372*** (0.0118)	0.0330** (0.0132)	0.0408*** (0.0130)
$Connected_{ij} \times \log(GDP_i)$	-0.0087 (0.0059)	-0.0086 (0.0057)	-0.0110* (0.0066)	-0.0105 (0.0064)
$Connected_{ij} \times \log(GDP_j)$	-0.0089* (0.0048)	-0.0093** (0.0047)	-0.0110* (0.0058)	-0.0108* (0.0057)
$Exposed_{ij}$		0.0410*** (0.0124)		0.0504*** (0.0142)
Observations	303,324	303,324	303,324	303,324
Squared Correlation	0.99297	0.99299	0.99331	0.99335
Pseudo R ²	0.99848	0.99849	0.99849	0.99850
BIC	1.01×10^{10}	1.01×10^{10}	1.01×10^{10}	1.01×10^{10}
Pair FE	✓	✓	✓	✓
Exporter-year FE	✓	✓	✓	✓
Importer-year FE	✓	✓	✓	✓

Notes: The coefficient on $Connected_{ij}$ reports the estimated treatment effect. In Columns (1)–(2), treatment is defined by a binary indicator equal to one if the country pair is directly connected by a cable and zero otherwise. In Columns (3)–(4), the treatment effect is estimated using cohort-by-time interactions. The coefficient on $Exposed_{ij}$ captures the exposure effect and is estimated using a binary indicator equal to one if the pair is classified as exposed and zero otherwise. $\log(Dist_{ij})$, $\log(GDP_i)$, and $\log(GDP_j)$ denote, respectively, the log bilateral distance, the log GDP of the exporter, and the log GDP of the importer, where log GDP is averaged over the five pre-treatment years. In Columns (1)–(2), these covariates are demeaned at the sample level; in Columns (3)–(4), they are demeaned within cohort-by-time cells. A Regional Trade Agreement dummy is also included as control. Robust standard errors clustered by pair in parentheses.

* = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

5.2 Dynamic and cohort-level effects

I now present the results from the event–study specification. Figure 3 displays the estimated dynamic effects from the main model, both with and without accounting for spillovers. The first point to note is that pre-treatment coefficients are not statistically different from zero, which is reassuring for the validity of both the parallel trends and no-anticipation assumptions. Post-treatment coefficients on the main effect are larger when spillovers are included, although the difference is not statistically significant. In the specification that accounts for spillovers, the average effect peaks at about 7.4% one year after the cable becomes ready for service. The positive effect then gradually declines but remains statistically significant up to four years after treatment. When spillovers are ignored, there is still a statistically significant increase in the first post-treatment year (about 4.4%), but the coefficients quickly lose significance from the second year onward.

Turning to the interaction terms, the distance interaction exhibits a clear positive jump at the RFS date, both with and without spillovers. In this case, accounting for spillovers yields only marginally larger post-treatment coefficients which, however, remain statistically significant up to six years after treatment. Lastly, the interactions with exporter and importer GDP exhibit a negative change at the treatment date, but most of the coefficients are imprecisely estimated.

Overall, there are at least two main considerations emerging from these results. The first is that the main event-study coefficient associated with the treatment year (i.e., year = 0) itself is relatively small compared to the peak in the second year. This is fully consistent with the nature of the treatment. In fact, the RFS date is defined as the year in which the cable is ready for service, but not all pairs benefit from the cable in that exact year. First, some cables enter service late (or very early) in the calendar year, so the first “treated” year may reflect only a few months of effective operation. Second, while the RFS date indicates when capacity becomes available, traffic is typically rerouted and adopted gradually, so the

economic effects are more likely to emerge in the subsequent years rather than immediately upon activation.

A second clear pattern is that the estimated effect fades over time. While this is interesting in its own right, it must be interpreted with care. Within the context of staggered event study designs the dynamic effects are not solely the result of the treatment, but they are also affected by the change in the cohort composition. In fact, the observations that remain at longer horizons come from cohorts that are treated relatively early in the sample, while late-cohorts primarily contribute to the first years after treatment. If early adopters were less affected by subsea cables—perhaps because digital technologies were still less developed (while late adopters may, for instance, benefit more), then the apparent fade-out may partly reflect this compositional shift across cohorts rather than a genuine decline in the effect over time. On the other hand, given the rapid evolution of fiber-optic technologies and upgrades to cable systems, it is also plausible that older cables may become obsolete over time relative to newer infrastructure ([TeleGeography, 2026a](#)), which could instead contribute to a real attenuation of the effect at longer horizons.

To address this concern, it is useful to examine the cohort-level treatment effects reported in Figure 4. While treatment effects vary across cohorts, they are remarkably stable overall, as most cohorts display estimates in the 5–8% range.¹² This pattern suggests that the apparent fade-out in the event-study coefficients is more plausibly driven by technological obsolescence rather than compositional shifts across cohorts.

Moreover, as for the estimates analyzed so far, the cohort results also make clear that accounting for spillovers tends to raise the estimated treatment effect. The size of this adjustment, however, is quite heterogeneous. For some cohorts, the gap between spillover-adjusted and unadjusted estimates is substantial, whereas for others (notably smaller cohorts such as 2007 and 2011) the two estimates almost coincide. A natural interpretation is that

¹²The few cohorts with small, zero, or negative coefficients are also among the least populated cohorts, as indicated with the distribution of new cable routes in Figure 1 (in the spillover-adjusted specification).

periods of intensive cable deployment generate larger spillovers that contaminate the control group and lead to a more severe understatement of the direct connection effect when spillovers are ignored. Conversely, when deployment activity is limited, spillovers likely affect fewer country pairs, so control-group contamination is weaker, and the adjusted and unadjusted estimates tend to coincide.

5.3 Testing different exposure cutoffs

As discussed in Section 4.1, to classify a pair as exposed or unexposed, the baseline specification uses a cutoff of 4 hops. That means that, if the shortest network distance between any two countries falls below (or becomes equal to) 4, then they are considered as exposed to spillovers. Given that the choice of the cutoff is mainly arbitrary, I show that the results are invariant to alternative cutoff specifications. In particular, I test for cutoffs of 3, 5, and 6 and, as expected, the main takeaway holds: controlling for spillovers makes the estimated treatment effect higher, regardless of the cutoff rule used. However, a few more notable aspects emerge from these additional estimates. In particular, the post-treatment estimates increase when going from no spillovers to a cutoff of 3 and then to the baseline cutoff of 4. The slightly smaller point estimate under the cutoff of 3 likely reflects the fact that this exposure rule is less restrictive and classifies fewer control units as exposed. Consequently, some pairs that are in fact affected by spillovers remain in the control group, slightly inflating the imputed control mean and attenuating the estimated treatment effect. In addition, the estimates show substantial convergence once higher cutoffs (5 and 6) are considered, as the coefficients are basically identical to baseline with cutoff 4. Overall, the fact that the coefficients are very similar under each of the alternative specifications indicates that the results are not sensitive to the precise choice of the cutoff once it is sufficiently inclusive, and it reinforces the robustness of the baseline choice of 4, which does not appear to drive the estimated coefficients.

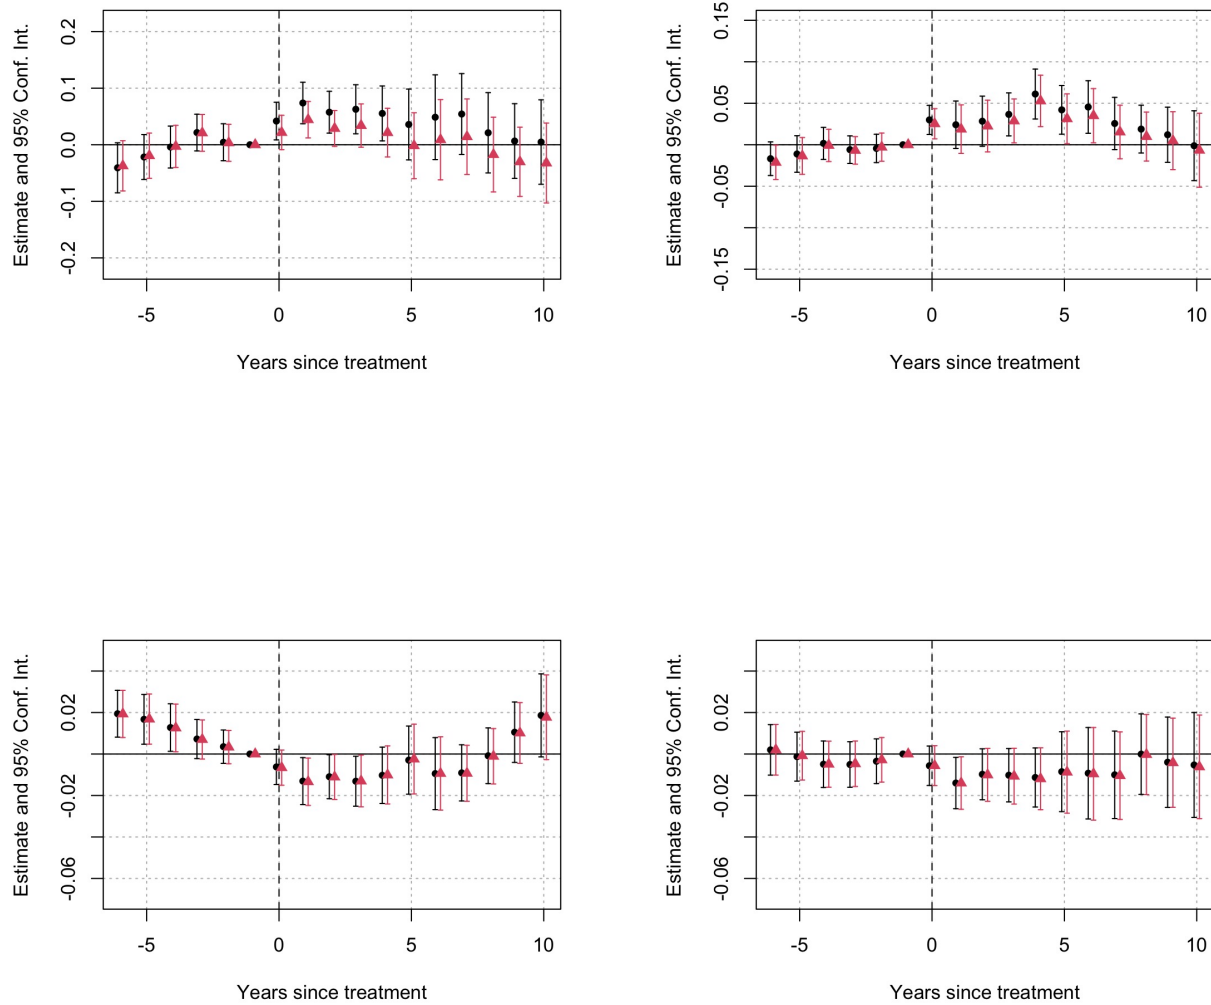


Figure 3: Event study estimates

Notes: The upper-left panel reports the event-study coefficients for the main treatment effect; the upper-right panel reports the interaction with distance; the lower-left and lower-right panels report the interactions with origin GDP and destination GDP, respectively. Black circles correspond to estimates that account for spillovers; red squares correspond to estimates that do not account for spillovers. Estimates that account for spillovers include interactions between $Exposed_{ij}$ and year dummies, allowing the exposure effect to vary over time. All specifications also control for the presence of a Regional Trade Agreement.




Figure 4: Cohort-level treatment effects

5.4 Previous exposure and treatment effects

Assumption 3 posits that, once a pair is treated, its subsequent outcomes no longer depend on its exposure history. As discussed in Section 4.1, this has two main implications. First, the model assumes that treated pairs do not receive additional spillovers after treatment, which is plausible in this setting. Second, treatment effects are assumed to be the same regardless of whether the pair was exposed to improved indirect routing before treatment, which is a more demanding restriction and thus needs to be further evaluated.

To study this new margin of heterogeneity, I relax the second implication by allowing the treatment effect to vary with prior exposure. I proceed in two steps. First, I redefine the exposure indicator so that it remains equal to one after treatment for units that were exposed before receiving a direct cable¹³. Second, I interact this new exposure indicator with the cohort-by-time treatment dummies, so that the post-treatment effect can differ between pairs that were exposed prior to treatment and those that were not.

The results support the presence of heterogeneity. The interaction between treatment and

¹³Mathematically: $e_{ijt} = \mathbf{1}\{t \geq EXP_{ij} \mid EXP_{ij} < RFS_{ij}\}$

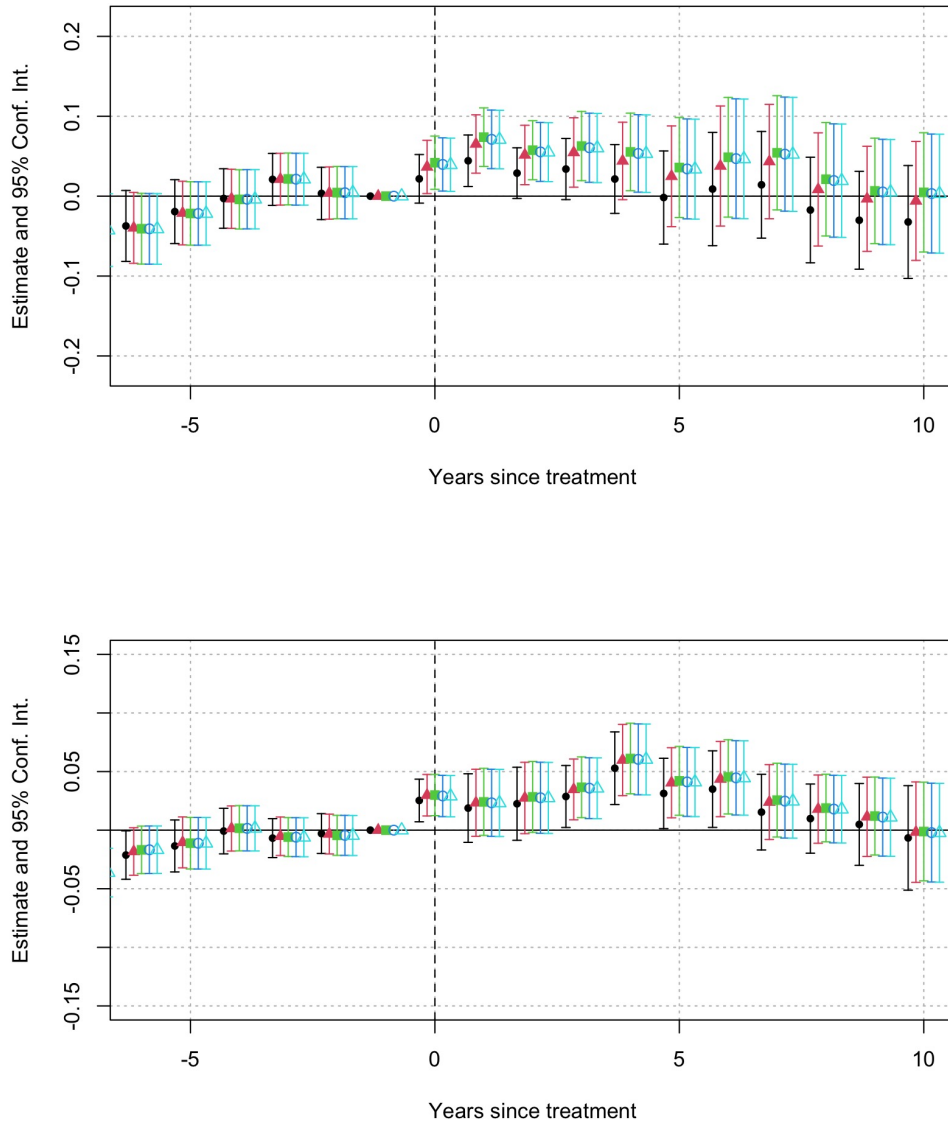


Figure 5: Event study estimates (Different cutoffs spec.)

Notes: The upper panel reports the event-study coefficients for the main treatment effect; the lower panel reports the interaction with distance. Black circles correspond to estimates that do not account for spillovers; red squares correspond to estimates that account for spillovers using three as distance cutoff; green squares correspond to estimates that account for spillovers using four as distance cutoff; blue hollow circles correspond to estimates that account for spillovers using five as distance cutoff; light blue hollow triangles correspond to estimates that account for spillovers using six as distance cutoff.

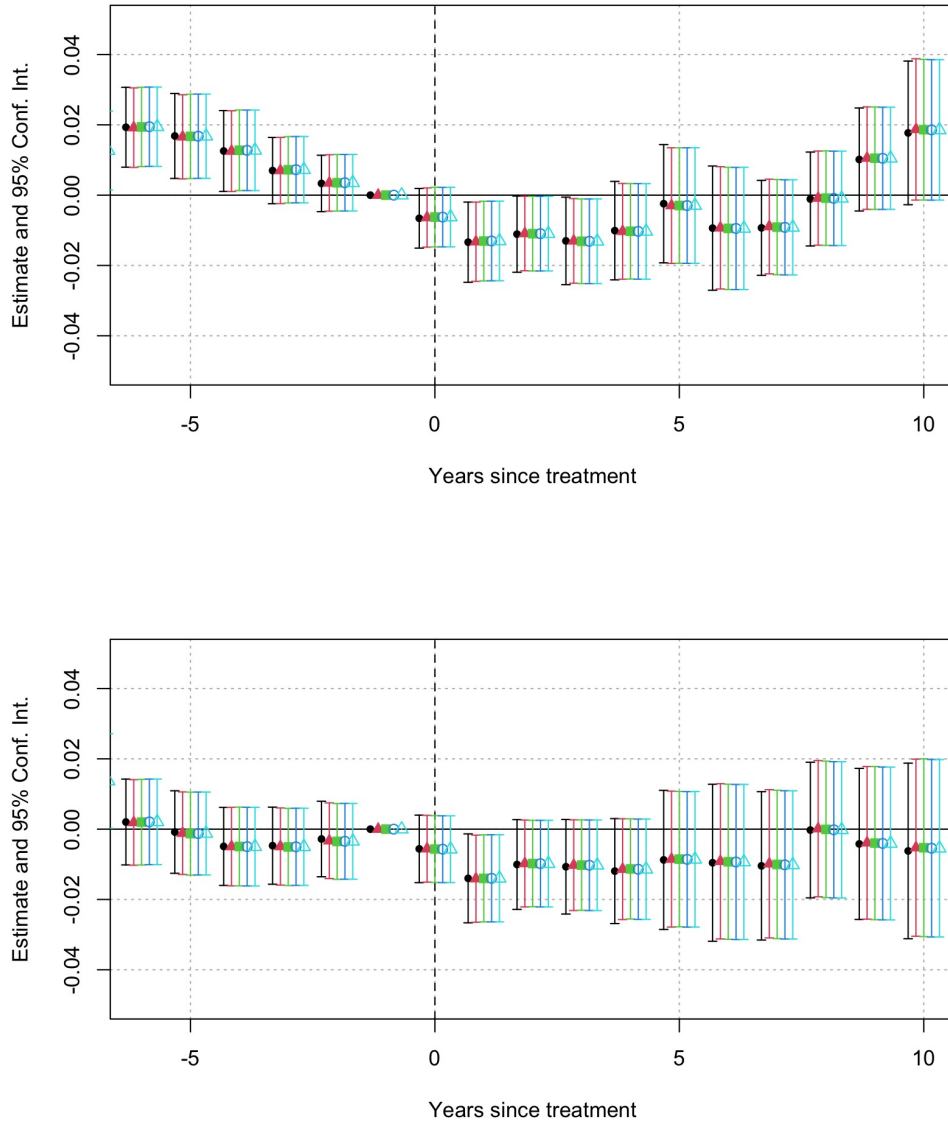


Figure 6: Event study estimates (Different cutoffs spec.)

Notes: the upper and lower panels report the interactions with origin GDP and destination GDP, respectively. Black circles correspond to estimates that do not account for spillovers; red squares correspond to estimates that account for spillovers using three as distance cutoff; green squares correspond to estimates that account for spillovers using four as distance cutoff; blue hollow circles correspond to estimates that account for spillovers using five as distance cutoff; light blue hollow triangles correspond to estimates that account for spillovers using six as distance cutoff. Estimates that account for spillovers include interactions between $Exposed_{ij}$ and year dummies, allowing the exposure effect to vary over time. All specifications also control for the presence of a Regional Trade Agreement.

prior exposure is negative, indicating that prior exposure attenuates the incremental effect of a new direct cable. Correspondingly, the main treatment effect (i.e., the coefficient associated with “ $Connected_{ij}$ ”) in this specification represents the implied estimate for treated pairs that were not exposed beforehand, and its magnitude increases relative to the pooled baseline estimate in column (1). Quantitatively, the implied effect for previously exposed pairs is about 2.96%¹⁴, whereas the implied effect for non-exposed pairs is about 7.71%¹⁵. This suggests that, while the effects are likely to be positive for both groups, previous exposure represents an important margin of heterogeneity, as countries that did not experience an improvement in indirect connectivity seem to benefit more from a new direct cable, as compared to countries that already enjoyed better indirect routes.

6 Conclusion

This paper studies how improvements in *bilateral* digital connectivity affect international trade. Submarine telecommunications cables, which carry the vast majority of international data traffic, constitute the core component of the infrastructure sustaining such connectivity and provide a natural setting to analyze how reductions in digital frictions translate into trade outcomes.

Empirically, I exploit the staggered timing of Ready-for-Service (RFS) dates for new direct cable links to implement a difference-in-differences design within a gravity framework. The analysis explicitly accounts for treatment spillovers by constructing an exposure mapping that accounts for country pairs that may be positively affected by the newly built cables. The main results indicate that establishing a direct cable connection increases bilateral trade and that not accounting for spillovers underestimates this effect. Dynamic estimates show that the effects emerge after the RFS year and remain positive for several subsequent periods.

¹⁴This estimate can be obtained by summing up the coefficients of $Connected_{ij}$, $Pot.Exposed_{ij}$ and their interaction.

¹⁵Previously exposed pairs account for roughly 53% of treated observations. The weighted average of 2.96% and 7.71% is 5.45%, very close to the baseline estimate where both groups are pooled.

Table 3: Aggregate treatment effects

	Dependent variable: $Trade_{ij}$	
	Robust TWFE (1)	Robust TWFE (2)
$Connected_{ij}$	0.0612*** (0.0205)	0.0771*** (0.0213)
$Exposed_{ij}$	0.0502*** (0.0141)	0.0360** (0.0162)
$Connected_{ij} \times \log(Dist_{ij})$	0.0404*** (0.0130)	0.0452*** (0.0131)
$Connected_{ij} \times \log(GDP_i)$	-0.0115* (0.0066)	-0.0110* (0.0065)
$Connected_{ij} \times \log(GDP_j)$	-0.0117** (0.0058)	-0.0115** (0.0058)
$Connected_{ij} \times Exposed_{ij}$		-0.0835*** (0.0210)
Observations	303,324	303,324
Squared Correlation	0.99335	0.99332
Pseudo R ²	0.99850	0.99850
BIC	1.01×10^{10}	1.01×10^{10}
Pair FE	✓	✓
Exporter-year FE	✓	✓
Importer-year FE	✓	✓

Notes: The coefficient on $Connected_{ij}$ reports the estimated treatment effect. The treatment effect is estimated using cohort-by-time interactions. The coefficient on $Exposed_{ij}$ captures the exposure effect and is estimated using a binary indicator equal to one if the pair is classified as exposed and zero otherwise. $\log(Dist_{ij})$, $\log(GDP_i)$, and $\log(GDP_j)$ denote, respectively, the log bilateral distance, the log GDP of the exporter, and the log GDP of the importer, where log GDP is averaged over the five pre-treatment years. Covariates demeaned within cohort-by-time cells. A Regional Trade Agreement dummy is also included as control. Robust standard errors clustered by pair in parentheses. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Treatment-effect heterogeneity further suggests that gains are larger for more distant partners, highlighting how improved connectivity reduces the role of distance in trade.

At the same time, the analysis features several limitations, some of which are data-related. A first constraint is that the publicly available dataset does not contain data on cables that have been dismantled before 2022, which limits the possibility to extend the analysis to earlier years. Moreover, the dataset misses key information on technical dimensions such as installed and lit capacity, fibre pairs, or upgrades.

A second limitation concerns the lack of data on terrestrial cables. Because inland and landlocked countries often rely on terrestrial backhaul and cross-border fibre to access submarine capacity, the absence of detailed terrestrial network data may lead to exposure misclassification and to an incomplete characterization of spillovers. Extending the network representation to jointly model subsea and terrestrial links would sharpen the interpretation of indirect exposure, but also enable a more policy-relevant assessment of infrastructure complementarities.

Another limitation concerns the specific construction of the exposure mapping. In particular, the chosen mapping does not allow testing spillovers operating on treated pairs after direct connection, because the shortest distance reaches the lower bound mechanically when pairs are treated. Future extensions could therefore aim to develop and implement exposure mappings that allow for variation in the exposure even post-treatment (e.g., through redundancy, congestion relief, or alternative-route improvements), allowing the estimation of marginal gains from subsequent network changes.

Another natural extension of this research is to move beyond aggregate bilateral goods trade and study which types of goods respond most strongly to improvements in connectivity. In particular, the effects of lower latency may be concentrated in digitally ordered, physically delivered transactions (e.g., cross-border parcel trade mediated by e-commerce platforms) or service-intensive manufactures whose sale and operation rely on complementary pre- and post-sale services. Focusing on these more connectivity-sensitive segments would help link the

reduced-form trade effects to the mechanisms discussed in Section 3.2 and provide sharper evidence on how digital infrastructure translates into merchandise trade outcomes.

Another possible direction is to move to firm-level outcomes. If increased connectivity reduces information and coordination frictions, their effects should be visible in micro-level margins such as entry into exporting and importing. Combining the infrastructure shock with customs or firm-panel data would allow the paper to distinguish whether aggregate gains are driven by more firms trading (extensive margin), incumbent firms expanding volumes (intensive margin), but also how these gains are distributed across different sectors.

Overall, the evidence presented in this paper supports the view that better digital infrastructure can stimulate trade between countries. By treating subsea cable deployments as pair-specific connectivity shocks and explicitly accounting for network spillovers, the paper provides a global, infrastructure-aware perspective on how the physical layer of the internet shapes real economic exchange.

References

- 3GPP (2023) “Specification 22.104: Service requirements for cyber-physical control applications in vertical domains,” 3GPP Specification Details (portal), <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3528>, Accessed 2026-03-01.
- Ahmed, Adnan, Zubair Shafiq, Harkeerat Bedi, and Amir Khakpour (2017) “Peering vs. Transit: Performance Comparison of Peering and Transit Interconnections,” in *2017 IEEE 25th International Conference on Network Protocols (ICNP)*, 1–10: IEEE, [10.1109/ICNP.2017.8117549](https://doi.org/10.1109/ICNP.2017.8117549).
- Akerman, Anders, Edwin Leuven, and Magne Mogstad (2022) “Information Frictions, Internet, and the Relationship between Distance and Trade,” *American Economic Journal: Applied Economics*, 14 (1), 133–163, [10.1257/app.20190589](https://doi.org/10.1257/app.20190589).
- Basalla, Marcus, Johannes Schneider, Martin Luksik, Roope Jaakonmäki, and Jan vom Brocke (2021) “On Latency of E-Commerce Platforms,” *Journal of Organizational Computing and Electronic Commerce*, 31 (1), 1–17, [10.1080/10919392.2021.1882240](https://doi.org/10.1080/10919392.2021.1882240).
- Bellucci, Chiara, Stela Rubínová, and Roberta Piermartini (2025) “Better Together: How Digital Connectivity and Regulation Reduce Trade Costs,” *Review of International Economics*, 33 (3), 796–814, [10.1111/roie.12807](https://doi.org/10.1111/roie.12807).
- Briglauer, Wolfgang, Carlo Cambini, Klaus Gugler, and Lorien Sabatino (2025) “Economic benefits of new broadband network coverage and service adoption: evidence from OECD member states,” *Industrial and Corporate Change*, 34 (4), 696–721, [10.1093/icc/dtae043](https://doi.org/10.1093/icc/dtae043).
- Brutlag, Jake (2009) “Speed Matters for Google Web Search,” Technical report, Google, Inc. https://services.google.com/fh/files/blogs/google_delayexp.pdf.

- Butts, Kyle (2023) “Difference-in-Differences Estimation with Spatial Spillovers,” [10.48550/arXiv.2105.03737](https://arxiv.org/abs/2105.03737), arXiv:2105.03737, version 3.
- Cadestin, Charles and Sébastien Miroudot (2020) “Services exported together with goods,” OECD Trade Policy Papers 236, OECD Publishing, Paris, [10.1787/275e520a-en](https://doi.org/10.1787/275e520a-en).
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant’Anna (2024) “Difference-in-Differences with a Continuous Treatment,” NBER Working Paper 32117, National Bureau of Economic Research, [10.3386/w32117](https://doi.org/10.3386/w32117).
- Callaway, Brantly and Pedro H. C. Sant’Anna (2021) “Difference-in-Differences with multiple time periods,” *Journal of Econometrics*, 225 (2), 200–230, [10.1016/j.jeconom.2020.12.001](https://doi.org/10.1016/j.jeconom.2020.12.001).
- Caoui, El Hadi and Andrew Steck (2025) “Content providers and the deployment of Internet infrastructure,” *Information Economics and Policy*, 71, 101134, [10.1016/j.infoecopol.2025.101134](https://doi.org/10.1016/j.infoecopol.2025.101134).
- Carballo, Jerónimo, Marisol Rodríguez Chatruc, Catalina Salas Santa, and Christian Volpe Martincus (2022) “Online Business Platforms and International Trade,” *Journal of International Economics*, 137, 103599, [10.1016/j.jinteco.2022.103599](https://doi.org/10.1016/j.jinteco.2022.103599).
- Cariolle, Joël (2021) “International connectivity and the digital divide in Sub-Saharan Africa,” *Information Economics and Policy*, 55, 100901, <https://doi.org/10.1016/j.infoecopol.2020.100901>.
- Casalini, Francesca and Javier López González (2019) “Trade and Cross-Border Data Flows,” OECD Trade Policy Papers 220, OECD Publishing, Paris, [10.1787/b2023a47-en](https://doi.org/10.1787/b2023a47-en).
- de Chaisemartin, Clément and Xavier D’Haultfœuille (2020) “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 110 (9), 2964–96, [10.1257/aer.20181169](https://doi.org/10.1257/aer.20181169).

Chang, Qing, Lin William Cong, Liyong Wang, and Longtian Zhang (2023) “Production, Trade, and Cross-Border Data Flows,” NBER Working Paper 31416, National Bureau of Economic Research, [10.3386/w31416](https://doi.org/10.3386/w31416), Revised March 2025.

Chang, Rocky K. C., Edmond W. W. Chan, Weichao Li, Waiting W. T. Fok, and Xiapu Luo (2010) “Could Ash Cloud or Deep-Sea Current Overwhelm the Internet?,” USENIX Hot Topics in System Dependability (HotDep), poster, October, https://www.usenix.org/event/hotdep10/tech/full_papers/Chang.pdf.

Clark, David, Steven Bauer, K. C. Claffy, Amogh Dhamdhere, Bradley Huffaker, William Lehr, and Matthew Luckie (2014) “Measurement and Analysis of Internet Interconnection and Congestion,” in *Telecommunications Policy Research Conference (TPRC)*, September, https://www-old.caida.org/publications/papers/2014/measurement_analysis_internet_interconnection/.

Conte, Maddalena, Pierre Cotterlaz, and Thierry Mayer (2022) “The CEPII Gravity Database,” Working Papers 2022-05, CEPII research center, <https://ideas.repec.org/p/cii/cepiddt/2022-05.html>.

Deloitte Ireland LLP and Fifty-Five (2020) “Milliseconds Make Millions: A study on how improvements in mobile site speed positively affect a brand’s bottom line,” Technical report, Commissioned by Google; prepared by Deloitte Ireland LLP based on data provided by Fifty-Five, https://www.thinkwithgoogle.com/_qs/documents/9757/Milliseconds_Make_Millions_report_hQYAbZJ.pdf, Mobile site speed monitored end-2019; report published 2020.

Demir, Banu, Beata Javorcik, and Piyush Panigrahi (2024) “Breaking Invisible Barriers: Does Fast Internet Improve Access to Input Markets?” December, https://piyush-panigrahi.com/FibreInternet_DJP.pdf, Working Paper.

- Deutschmann, Jörg, Kai-Steffen Hielscher, and Reinhard German (2019) “Satellite Internet Performance Measurements,” in *Proceedings of the 2019 International Conference on Networked Systems (NetSys 2019)*: IEEE, [10.1109/NetSys.2019.8854494](https://doi.org/10.1109/NetSys.2019.8854494).
- D’Andrea, Angelo and Nicola Limodio (2024) “High-Speed Internet, Financial Technology, and Banking,” *Management Science*, 70 (2), 773–798, [10.1287/mnsc.2023.4703](https://doi.org/10.1287/mnsc.2023.4703).
- ETSI (2022) “ETSI TS 122 104 V17.7.0: 5G; Service requirements for cyber-physical control applications in vertical domains (3GPP TS 22.104 Release 17),” Technical Specification TS 122 104 V17.7.0, European Telecommunications Standards Institute, https://www.etsi.org/deliver/etsi_ts/122100_122199/122104/17.07.00_60/ts_122104v170700p.pdf.
- (2024) “ETSI TR 103 702 V1.2.1: Speech and multimedia Transmission Quality (STQ); QoS parameters and test scenarios for assessing network capabilities in 5G performance measurements,” Technical Report TR 103 702 V1.2.1, European Telecommunications Standards Institute, https://www.etsi.org/deliver/etsi_tr/103700_103799/103702/01.02.01_60/tr_103702v010201p.pdf.
- Fang, Chao, Xuepeng Tang, Linhui Yu, and Junsen Zhang (2024) “Connecting to Trade: Submarine Optical Cables and Chinese Firms’ Exports,” <https://ssrn.com/abstract=5254120>, Preprint (not peer reviewed).
- Fanou, Rodéric, Bradley Huffaker, Ricky Mok, and K. C. Claffy (2020) “Unintended Consequences: Effects of Submarine Cable Deployment on Internet Routing,” in *Passive and Active Measurement*, 12048 of Lecture Notes in Computer Science, 211–227: Springer, [10.1007/978-3-030-44081-7_13](https://doi.org/10.1007/978-3-030-44081-7_13).
- Fernandes, Ana M., Aaditya Mattoo, Huy Nguyen, and Marc Schiffbauer (2019) “The Internet and Chinese Exports in the Pre-Alibaba Era,” *Journal of Development Economics*, 138, 57–76, [10.1016/j.jdeveco.2018.11.003](https://doi.org/10.1016/j.jdeveco.2018.11.003).

- Ferracane, Martina F., Bernard Hoekman, Erik van der Marel, and Filippo Santi (2023) “Digital Trade, Data Protection and EU Adequacy Decisions,” CITP Working Paper 006, Centre for Inclusive Trade Policy (CITP), https://citp.ac.uk/wp-content/uploads/Ferracane-et-al-Adequacy-paper_Oct2023_CITP_WP6.pdf, Published 11 October 2023.
- Fiorini, Mario, Wooyong Lee, and Gregor Pfeifer (2024) “A Simple Approach to Staggered Difference-in-Differences in the Presence of Spillovers,” IZA Discussion Paper 16868, IZA Institute of Labor Economics, <https://docs.iza.org/dp16868.pdf>, IZA Discussion Paper No. 16868.
- Freund, Caroline L. and Diana Weinhold (2004) “The Effect of the Internet on International Trade,” *Journal of International Economics*, 62 (1), 171–189, [10.1016/S0022-1996\(03\)00059-X](https://doi.org/10.1016/S0022-1996(03)00059-X).
- Ganapati, Sharat and Woan Foong Wong (2023) “How Far Goods Travel: Global Transport and Supply Chains from 1965–2020,” *Journal of Economic Perspectives*, 37 (3), 3–30, [10.1257/jep.37.3.3](https://doi.org/10.1257/jep.37.3.3).
- Goodman-Bacon, Andrew (2021) “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 225 (2), 254–277, [10.1016/j.jeconom.2021.03.014](https://doi.org/10.1016/j.jeconom.2021.03.014).
- Head, Keith and John Ries (1999) “Rationalization effects of tariff reductions,” *Journal of International Economics*, 47 (2), 295–320, [https://doi.org/10.1016/S0022-1996\(98\)00019-1](https://doi.org/10.1016/S0022-1996(98)00019-1).
- Hjort, Jonas and Jonas Poulsen (2019) “The Arrival of Fast Internet and Employment in Africa,” *American Economic Review*, 109 (3), 1032–79, [10.1257/aer.20161385](https://doi.org/10.1257/aer.20161385).
- Hsu, Shih-Kai, James Kuo, Chung-Liang Lin, Ching-Hui Tsai, Wen-Bin Doo, Chin-Yi Ku, and Jean-Claude Sibuet (2008) “Turbidity currents, submarine landslides and the 2006 Pingtung earthquake off SW Taiwan,” *TAO: Terrestrial, Atmospheric and Oceanic Sciences*, 19 (6), 7–20, <https://tao.tcgu.org.tw/index.php/2008/200819-6>.

- Huffaker, Bradley, Marina Fomenkov, Daniel Plummer, David Moore, and K. C. Claffy (2002) “Distance Metrics in the Internet,” in *IEEE International Telecommunications Symposium (ITS)*, 200–202, September, <https://www-old.caida.org/publications/papers/2002/Distance/>.
- Imbruno, Michele, Joël Cariolle, and Jaime de Melo (2025) “Digital connectivity and firm participation in foreign markets: An exporter-based bilateral analysis,” *Journal of Development Economics*, 177, 103551, [10.1016/j.jdeveco.2025.103551](https://doi.org/10.1016/j.jdeveco.2025.103551).
- International Telecommunication Union (2001) “ITU-T Recommendation G.1010: End-user multimedia QoS categories,” Recommendation G.1010, ITU-T, <https://www.itu.int/rec/T-REC-G.1010-200111-I>.
- (2003) “ITU-T Recommendation G.114: One-way transmission time,” Recommendation G.114, ITU-T, <https://www.itu.int/rec/T-REC-G.114>, In-force version (05/2003) with subsequent incorporated amendments.
- Lendle, Andreas, Marcelo Olarreaga, Simon Schropp, and Pierre-Louis Vézina (2016) “There Goes Gravity: eBay and the Death of Distance,” *The Economic Journal*, 126 (591), 406–441, <https://doi.org/10.1111/eoj.12286>.
- Liu, Shucheng, Zachary S. Bischof, Ishaan Madan, Peter K. Chan, and Fabián E. Bustamante (2020) “Out of Sight, Not Out of Mind: A User-View on the Criticality of the Submarine Cable Network,” in *Proceedings of the ACM Internet Measurement Conference (IMC '20)*: ACM, [10.1145/3419394.3423633](https://doi.org/10.1145/3419394.3423633).
- Livadariu, Ioana, Ahmed Elmokashfi, and Georgios Smaragdakis (2024) “Tracking Submarine Cables in the Wild,” *Computer Networks*, 242, 110234, [10.1016/j.comnet.2024.110234](https://doi.org/10.1016/j.comnet.2024.110234).
- López González, Javier and Janos Ferencz (2018) “Digital Trade and Market Openness,” Technical Report 217, OECD Publishing, Paris, [10.1787/1bd89c9a-en](https://doi.org/10.1787/1bd89c9a-en).

- López González, Javier and Silvia Sorescu (2021) “Trade in the Time of Parcels,” OECD Trade Policy Papers 249, OECD Publishing, Paris, [10.1787/0faac348-en](https://doi.org/10.1787/0faac348-en).
- Malgouyres, Clément, Thierry Mayer, and Clément Mazet-Sonilhac (2021) “Technology-induced trade shocks? Evidence from broadband expansion in France,” *Journal of International Economics*, 133, 103520, <https://doi.org/10.1016/j.jinteco.2021.103520>.
- Mauldin, Alan (2023) “Do Submarine Cables Account For Over 99% of Intercontinental Data Traffic?,” <https://blog.telegeography.com/2023-mythbusting-part-3>, TeleGeography infrastructure mythbusting.
- Miroudot, Sébastien (2017) “The Servicification of Global Value Chains: Evidence and Policy Implications,” paper prepared for the unctad multi-year expert meeting on trade, services and development (5th session), UNCTAD, Geneva, https://unctad.org/system/files/non-official-document/c1mem5_2017_124_S3_Miroudot_2.pdf.
- Miroudot, Sébastien and Charles Cadestin (2017) “Services in Global Value Chains: From Inputs to Value-Creating Activities,” OECD Trade Policy Papers 197, OECD Publishing, Paris, [10.1787/465f0d8b-en](https://doi.org/10.1787/465f0d8b-en).
- Nagengast, Arne J. and Yoto V. Yotov (2025) “Staggered Difference-in-Differences in Gravity Settings: Revisiting the Effects of Trade Agreements,” *American Economic Journal: Applied Economics*, 17 (1), 271–96, [10.1257/app.20230089](https://doi.org/10.1257/app.20230089).
- Pope, Ed L., Peter J. Talling, Lionel Carter, Michael A. Clare, and James E. Hunt (2017) “Damaging sediment density flows triggered by tropical cyclones,” *Earth and Planetary Science Letters*, 458, 161–169, [10.1016/j.epsl.2016.10.046](https://doi.org/10.1016/j.epsl.2016.10.046).
- Romalis, John (2007) “NAFTA’s and CUSFTA’s Impact on International Trade,” *The Review of Economics and Statistics*, 89 (3), 416–435, [10.1162/rest.89.3.416](https://doi.org/10.1162/rest.89.3.416).

Spring, Neil, Ratul Mahajan, and Thomas Anderson (2003) “Quantifying the Causes of Path Inflation,” in *Proceedings of the ACM SIGCOMM 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, 113–124: ACM, [10.1145/863955.863970](https://doi.org/10.1145/863955.863970).

Stadnik, Wiktor and Ziemowit Nowak (2018) “The Impact of Web Pages’ Load Time on the Conversion Rate of an E-Commerce Platform,” in Borzemski, Leszek, Jerzy Świątek, and Zofia Wilimowska eds. *Information Systems Architecture and Technology: Proceedings of 38th International Conference on Information Systems Architecture and Technology – ISAT 2017*, 655 of *Advances in Intelligent Systems and Computing*, 336–345, Cham: Springer, [10.1007/978-3-319-67220-5_31](https://doi.org/10.1007/978-3-319-67220-5_31).

Steinwender, Claudia (2018) “Real Effects of Information Frictions: When the States and the Kingdom Became United,” *American Economic Review*, 108 (3), 657–696, [10.1257/aer.20150681](https://doi.org/10.1257/aer.20150681).

Straits Research (2024) “Submarine Cables Market Size, Share & Trends Analysis Report, 2033,” <https://straitsresearch.com/report/submarine-cables-market>.

Sun, Liyang and Sarah Abraham (2021) “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225 (2), 175–199, [10.1016/j.jeconom.2020.09.006](https://doi.org/10.1016/j.jeconom.2020.09.006).

TeleGeography (2025) “State of the Network 2025,” https://www2.telegeography.com/hubfs/LP-Assets/Ebooks/state-of-the-network-2025.pdf?_hsenc=p2ANqtz-8VZqTS78qeEana16KlFyqncs2DfioYnFHNVHBK_LoNBwLzMrOQY1VBmTAPTvzIPiiaGfM-Twugt1P2pKTLsfo0zLhS8MVkq7b2d2wTdUBoacqAaU&_hsmi=343529837, New submarine cable investment has averaged over USD 2 billion per year and total value of cables entering service 2024–2026 forecasted to exceed USD 10 billion.

TeleGeography (2026a) “Submarine Cable FAQs (Frequently Asked Questions),” <https://www2.telegeography.com/submarine-cable-faqs-frequently-asked-questions>, Accessed 2026-03-09.

——— (2026b) “Submarine Cable Map,” <https://www.submarinecablemap.com/>, Accessed: 2026-03-04.

UNCTAD (2024) “Review of Maritime Transport 2024, Chapter I: International maritime trade,” https://unctad.org/system/files/official-document/rmt2024ch1_en.pdf, Accessed: 2026-03-06.

——— (2025) “Meeting on Servicification and Trade in Services: Unlocking Opportunities for Developing Economies,” Joint MOFCOM–UNCTAD event at the China International Fair for Trade in Services (CIFTIS), Beijing, China, 12 September 2025, September, <https://unctad.org/meeting/meeting-servicification-and-trade-services-unlocking-opportunities-developing-economies>, Accessed: 2026-02-28.

United Nations (2017) “Submarine Cables and Pipelines,” in *The First Global Integrated Marine Assessment (World Ocean Assessment I)*: Cambridge University Press, [10.1017/9781108186148.022](https://doi.org/10.1017/9781108186148.022), Chapter 19, published online 18 May 2017.

Völk, Florian, Thomas Schlichter, Sumit Kumar et al. (2024) “5G Non-Terrestrial Networks With OpenAirInterface: An Experimental Study Over GEO Satellites,” *IEEE Access*, 12, 155098–155109, [10.1109/ACCESS.2024.3483958](https://doi.org/10.1109/ACCESS.2024.3483958).

Warf, Barney (2011) “Engineering Time and Space with the Global Fiber Optics Industry,” in Brunn, Stanley D. ed. *Engineering Earth*, 115–129, Dordrecht: Springer, [10.1007/978-90-481-9920-4_8](https://doi.org/10.1007/978-90-481-9920-4_8).

- Winseck, Dwayne (2017) “The Geopolitical Economy of the Global Internet Infrastructure,” *Journal of Information Policy*, 7, 228–267, [10.5325/jinfopoli.7.2017.0228](https://doi.org/10.5325/jinfopoli.7.2017.0228).
- Wooldridge, Jeffrey M. (2023) “Simple Approaches to Nonlinear Difference-in-Differences with Panel Data,” *The Econometrics Journal*, 26 (3), C31–C66, [10.1093/ectj/utad016](https://doi.org/10.1093/ectj/utad016).
- World Economic Forum (2021) “Regionalization vs globalization: future direction of trade?,” <https://www.weforum.org/stories/2021/07/regionalization-globalization-future-direction-trade/>, Accessed: 2026-03-06.
- Yotov, Yoto V., Roberta Piermartini, José-Antonio Monteiro, and Mario Larch (2016) *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*, Geneva: World Trade Organization and United Nations, https://www.wto.org/english/res_e/booksp_e/advancedwtounctad2016_e.pdf.

7 Appendix

7.1 Additional Estimation results

7.1.1 First cable specification

Below, I report results from the specification in which, for each country pair, I retain only the first cable built from 2005 onward and disregard any additional cables constructed subsequently. Treatment-effect estimates in Table 4 are slightly smaller than the main estimates reported in Table 2 and are less statistically significant. This is expected, as restricting the sample in this way reduces the number of treated observations and therefore lowers statistical power, making the estimates noisier. On the contrary, the treatment–distance interaction remains large and statistically significant across specifications.

The event-study estimates in Figure 7 are also broadly consistent with those obtained using the baseline sample. If anything, the event-study profiles appear less stable here than in the main specification, which further supports the use of the reweighting approach. Overall, these results suggest that weighting by the number of cables is not critical for the robustness of the findings.

7.1.2 Post-2000 Specification

Below, I report results from the baseline specification when I expand the treatment definition to include cables constructed from 2000 onward. As shown in Table 5, the estimated coefficients are broadly similar to those obtained in the baseline sample. If anything, the robust estimates that account for spillovers are slightly smaller in magnitude than in the main specification. The event-study results in Figure 8 make this difference even clearer: post-treatment effects are substantially weaker in the years following cable connection relative to the baseline specification.

Table 6 also reports estimates from an augmented specification that allows treatment

Table 4: Aggregate treatment effects (First cable spec.)

	Dependent variable: $Trade_{ij}$			
	TWFE	TWFE	Robust TWFE	Robust TWFE
$Connected_{ij}$	0.0102 (0.0113)	0.0247* (0.0135)	0.0079 (0.0139)	0.0304* (0.0164)
$Connected_{ij} \times \log(Dist_{ij})$	0.0217** (0.0088)	0.0269*** (0.0091)	0.0267*** (0.0099)	0.0280*** (0.0101)
$Connected_{ij} \times \log(GDP_i)$	-0.0048 (0.0038)	-0.0047 (0.0039)	-0.0052 (0.0048)	-0.0055 (0.0048)
$Connected_{ij} \times \log(GDP_j)$	-0.0019 (0.0032)	-0.0021 (0.0034)	-0.0020 (0.0048)	-0.0022 (0.0048)
$Exposed_{ij}$		0.0251** (0.0106)		0.0374*** (0.0138)
Observations	297,696	297,696	297,696	297,696
Squared Correlation	0.99365	0.99364	0.99423	0.99424
Pseudo R ²	0.99326	0.99326	0.99332	0.99332
BIC	1.15×10^{10}	1.15×10^{10}	1.14×10^{10}	1.14×10^{10}
Pair FE	✓	✓	✓	✓
Exporter-year FE	✓	✓	✓	✓
Importer-year FE	✓	✓	✓	✓

Notes: The coefficient on $Connected_{ij}$ reports the estimated treatment effect. In Columns (1)–(2), treatment is defined by a binary indicator equal to one if the country pair is directly connected by a cable and zero otherwise. In Columns (3)–(4), the treatment effect is estimated using cohort-by-time interactions. The coefficient on $Exposed_{ij}$ captures the exposure effect and is estimated using a binary indicator equal to one if the pair is classified as exposed and zero otherwise. $\log(Dist_{ij})$, $\log(GDP_i)$, and $\log(GDP_j)$ denote, respectively, the log bilateral distance, the log GDP of the exporter, and the log GDP of the importer, where log GDP is averaged over the five pre-treatment years. In Columns (1)–(2), these covariates are demeaned at the sample level; in Columns (3)–(4), they are demeaned within cohort-by-time cells. A Regional Trade Agreement dummy is also included as control. Robust standard errors clustered by pair in parentheses.

* = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

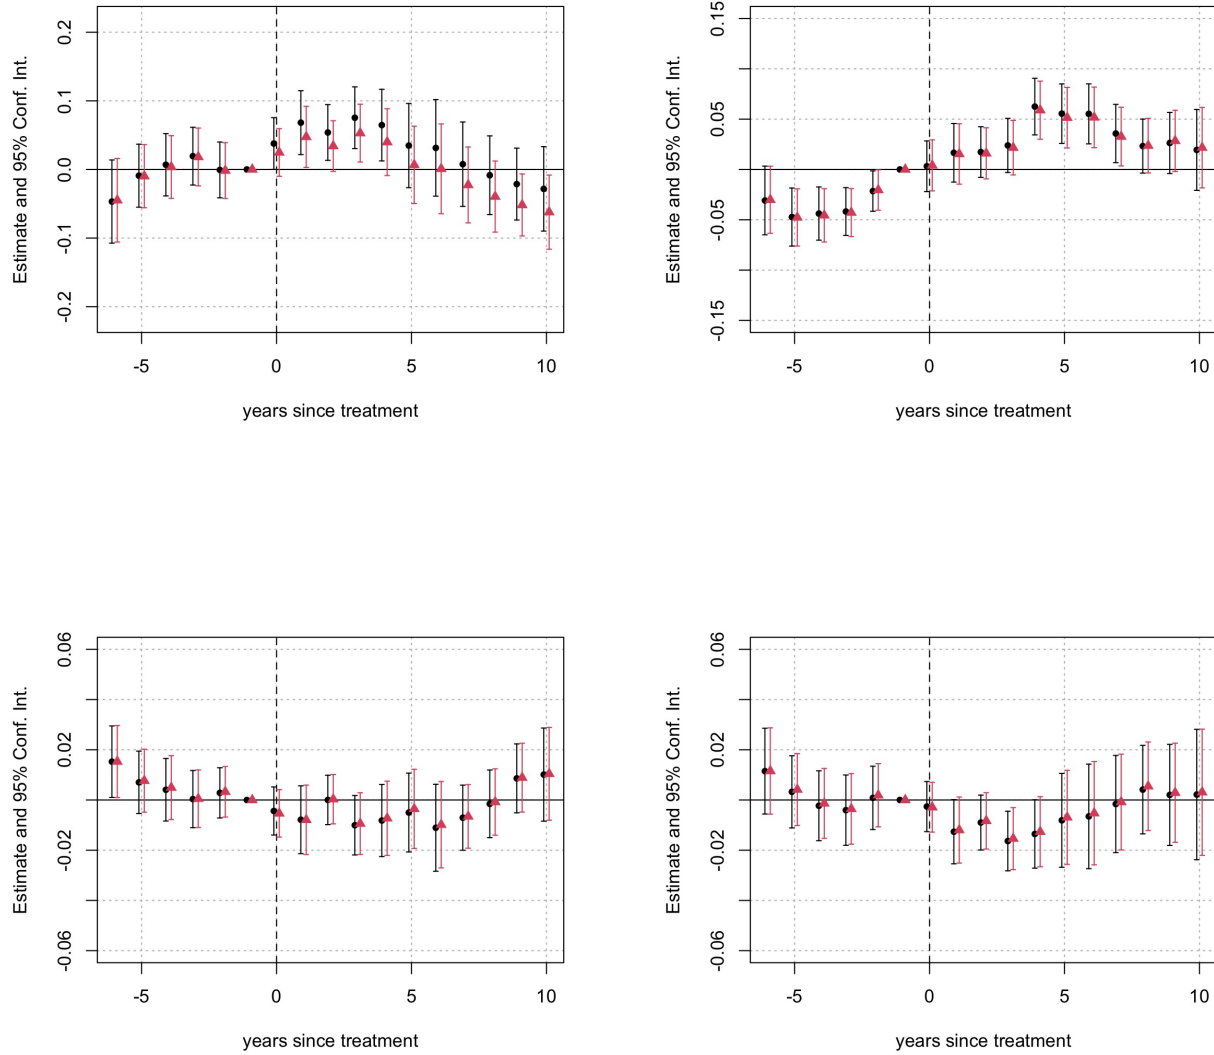


Figure 7: Event study estimates (First cable spec.)

Notes: The upper-left panel reports the event-study coefficients for the main treatment effect; the upper-right panel reports the interaction with distance; the lower-left and lower-right panels report the interactions with origin GDP and destination GDP, respectively. Black circles correspond to estimates that account for spillovers; red squares correspond to estimates that do not account for spillovers. Estimates that account for spillovers include interactions between $Exposed_{ij}$ and year dummies, allowing the exposure effect to vary over time. All specifications also control for the presence of a Regional Trade Agreement.

effects to vary with prior exposure. The results remain broadly consistent with the baseline specification but yield smaller effects. In particular, pairs that were not previously exposed experience a 5.6% increase in trade following cable construction, whereas the implied estimate for previously exposed pairs is close to zero (approximately 0.71%).

Overall, these findings indicate that including the full sample of cables does not qualitatively alter the main conclusions. However, the smaller estimated effects are consistent with the expectation that the early-2000s wave of cable construction had a more limited impact than later cables.

Table 5: Aggregate treatment effects (Post-2000 spec.)

	Dependent variable: $Trade_{ij}$			
	TWFE	TWFE	Robust TWFE	Robust TWFE
$Connected_{ij}$	0.0252* (0.0132)	0.0315*** (0.0083)	0.0217*** (0.0063)	0.0378*** (0.0085)
$Connected_{ij} \times \log(Dist_{ij})$	0.0237** (0.0100)	0.0198*** (0.0063)	0.0169** (0.0069)	0.0198*** (0.0070)
$Connected_{ij} \times \log(GDP_i)$	-0.0095** (0.0047)	-0.0055** (0.0026)	-0.0066*** (0.0025)	-0.0061** (0.0026)
$Connected_{ij} \times \log(GDP_j)$	-0.0072** (0.0037)	-0.0049** (0.0022)	-0.0054** (0.0022)	-0.0050** (0.0024)
$Exposed_{ij}$		0.0234*** (0.0079)		0.0307*** (0.0092)
Observations	312,291	312,291	312,291	312,291
Squared Correlation	0.99284	0.99484	0.99519	0.99519
Pseudo R ²	0.99865	0.99435	0.99439	0.99439
BIC	1.07×10^{10}	1.37×10^{10}	1.36×10^{10}	1.36×10^{10}
Pair FE	✓	✓	✓	✓
Exporter-year FE	✓	✓	✓	✓
Importer-year FE	✓	✓	✓	✓

Notes: The coefficient on $Connected_{ij}$ reports the estimated treatment effect. In Columns (1)–(2), treatment is defined by a binary indicator equal to one if the country pair is directly connected by a cable and zero otherwise. In Columns (3)–(4), the treatment effect is estimated using cohort-by-time interactions. The coefficient on $Exposed_{ij}$ captures the exposure effect and is estimated using a binary indicator equal to one if the pair is classified as exposed and zero otherwise. $\log(Dist_{ij})$, $\log(GDP_i)$, and $\log(GDP_j)$ denote, respectively, the log bilateral distance, the log GDP of the exporter, and the log GDP of the importer, where log GDP is averaged over the five pre-treatment years. In Columns (1)–(2), these covariates are demeaned at the sample level; in Columns (3)–(4), they are demeaned within cohort-by-time cells. A Regional Trade Agreement dummy is also included as control. Robust standard errors clustered by pair in parentheses.

* = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

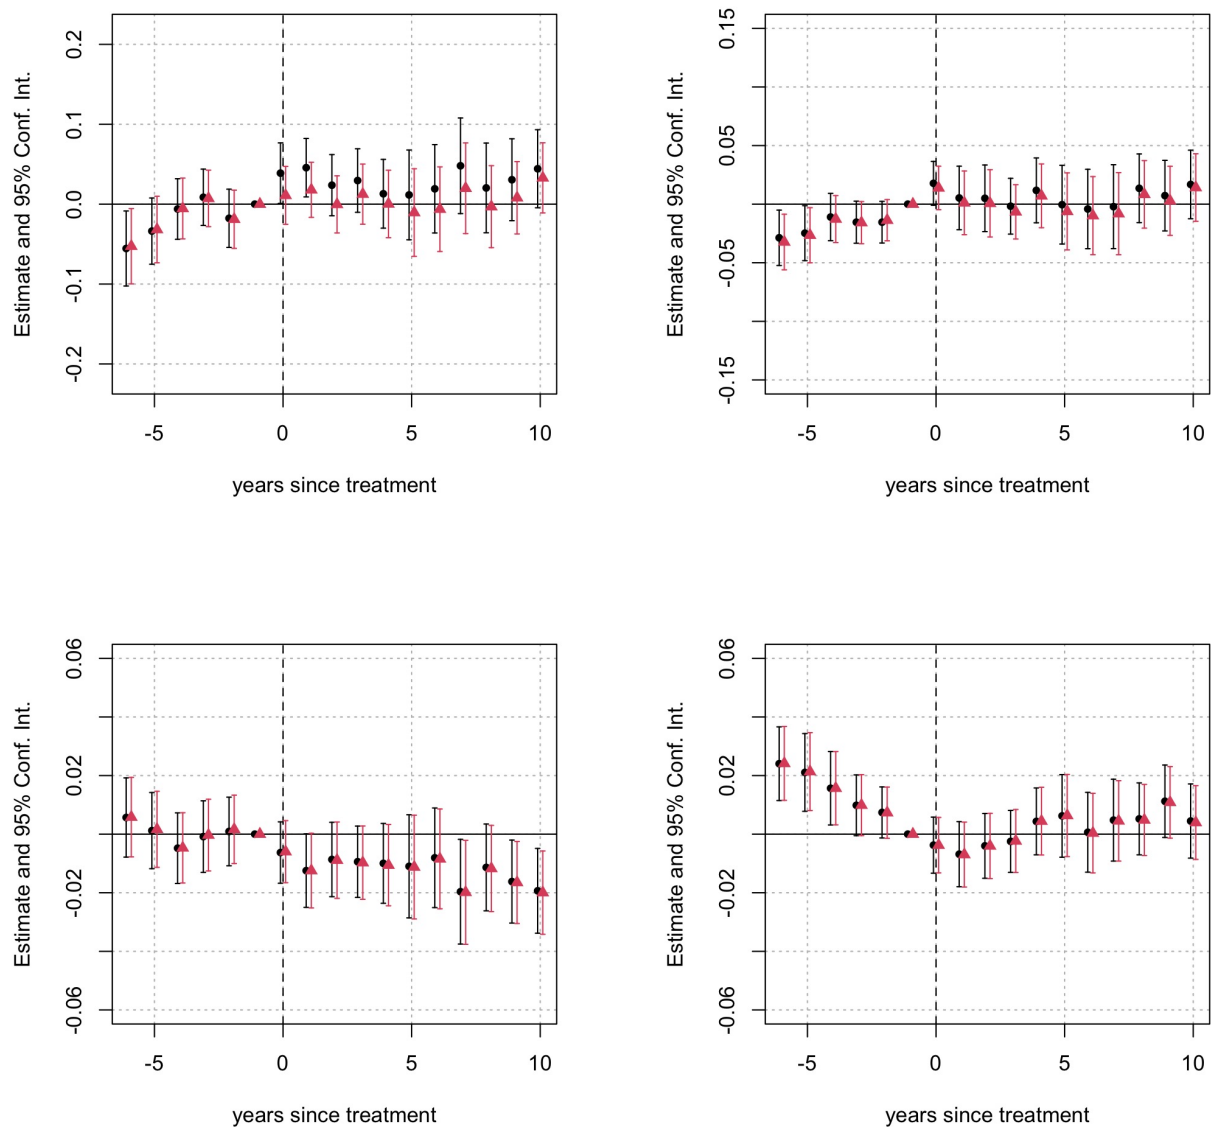


Figure 8: Event study estimates (Post-2000 spec.)

Notes: The upper-left panel reports the event-study coefficients for the main treatment effect; the upper-right panel reports the interaction with distance; the lower-left and lower-right panels report the interactions with origin GDP and destination GDP, respectively. Black circles correspond to estimates that account for spillovers; red squares correspond to estimates that do not account for spillovers. Estimates that account for spillovers include interactions between $Exposed_{ij}$ and year dummies, allowing the exposure effect to vary over time. All specifications also control for the presence of a Regional Trade Agreement.

Table 6: Aggregate treatment effects (Post-2000 spec.)

	Dependent variable: $Trade_{ij}$	
	Robust TWFE (1)	Robust TWFE (2)
$Connected_{ij}$	0.0378*** (0.0085)	0.0557*** (0.0142)
$Connected_{ij} \times \log(Dist_{ij})$	0.0198*** (0.0070)	0.0361*** (0.0108)
$Connected_{ij} \times \log(GDP_i)$	-0.0061** (0.0026)	-0.0094* (0.0049)
$Connected_{ij} \times \log(GDP_j)$	-0.0050** (0.0024)	-0.0067 (0.0042)
$Exposed_{ij}$	0.0307*** (0.0092)	0.0326** (0.0132)
$Connected_{ij} \times Exposed_{ij}$		-0.0812*** (0.0180)
Observations	312,291	312,291
Squared Correlation	0.99519	0.99327
Pseudo R ²	0.99439	0.99866
BIC	1.36×10^{10}	1.06×10^{10}
Pair FE	✓	✓
Exporter-year FE	✓	✓
Importer-year FE	✓	✓

Notes: The coefficient on $Connected_{ij}$ reports the estimated treatment effect. The treatment effect is estimated using cohort-by-time interactions. The coefficient on $Exposed_{ij}$ captures the exposure effect and is estimated using a binary indicator equal to one if the pair is classified as exposed and zero otherwise. $\log(Dist_{ij})$, $\log(GDP_i)$, and $\log(GDP_j)$ denote, respectively, the log bilateral distance, the log GDP of the exporter, and the log GDP of the importer, where log GDP is averaged over the five pre-treatment years. Covariates demeaned within cohort-by-time cells. A Regional Trade Agreement dummy is also included as control. Robust standard errors clustered by pair in parentheses. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.