Digital Interdependence and Power Politics
Supplementary Material

Last modified: March 13, 2025

Contents

1 Internet Interconnection Data

1.1 Overview and Collection

The Center for Applied Internet Data Analysis (CAIDA) at the University of California, San Diego has gathered data about different aspects of the internet architecture since 1998. This paper leverages two CAIDA datasets, *AS Relationships* with peering agreements between systems,[1] and *AS Organizations* that maps autonomous system (AS) numbers to organizations.[2] These independent operators form agreements to exchange data between one another through the Border Gateway Protocol (BGP). The internet needs a map to understand how data should get from one node to another. This protocol is how the internet routes requests for information between two systems.

The researchers at CAIDA collect routing data through snapshots of the BGP table over a five day period. These measurements are typically taken between the 1st and 5th of the month. The snapshots of the BGP table come from large collaborative internet infrastructure measurement projects including Route Views and the Réseaux IP Européens Network Coordination Centre (RIPE) Routing Information Service (RIS). These organizations collect data using a network of routers and collectors located at various points throughout the Internet. These routers are configured to collect BGP updates, which are messages sent between routers to announce changes in routing. The BGP updates collected by RouteViews are processed and stored in a database, which can be accessed by researchers and network operators. The data is used to study internet routing, identify routing problems and anomalies, and develop tools and techniques to improve Internet routing stability and security.

This dataset provides the set of efficient paths that data can take through the internet.[3] However, these pathways can either be peer-to-peer or customer-to-provider. The researchers from CAIDA develop an algorithm to infer the types of paths for each connection in the routing table. Several papers have validated CAIDA's approach to measuring BGP relationships.[4] However, it is possible that some peering relationships are wrongly coded as provider-to-customer relationships. We have no reason to believe that this error would differ over time or correlate with alliances or conflict.

[1] https://www.caida.org/data/as-relationships/

[2] https://www.caida.org/data/as-organizations/

[3] For some resent work on the continued importance of data transit and peering as content delivery networks proliferate see Huston (2016); Kolkman et al. (2022).

[4] Dimitropoulos, Xenofontas, Dmitri Krioukov, Marina Fomenkov, Bradley Huffaker, Young Hyun, kc claffy, and George Riley. "AS Relationships: Inference and Validation." ACM SIGCOMM Computer Communication Review 37, no. 1 (January 22, 2007): 29–40. https://doi.org/10.1145/1198255.1198259.; Luckie, Matthew, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and kc claffy. "AS Relationships, Customer Cones, and Validation." In Proceedings of the 2013 Conference on Internet Measurement Conference, 243–56. Barcelona Spain: ACM, 2013. https://doi.org/10.1145/2504730.2504735.; Dimitropoulos, Xenofontas, Dmitri Krioukov, Bradley Huffaker, kc claffy, and George Riley. "Inferring AS Relationships: Dead End or Lively Beginning?" In International Workshop on Experimental and Efficient Algorithms, 113–25, 2005. https://doi.org/10.1007/11427186_12.

1.2 Method for Creating Cross-sectional Autonomous System Ownership and Location Dataset

The most significant data collection and cleaning effort in this paper was to create a time-series-cross-sectional dataset of autonomous system ownership and location. This does not exist as one dataset, and there are several challenges to creating one. CAIDA's AS Organizations dataset from 2004 contains results from quarterly bulk dumps (usually the first day of January, April, July and October) from the WHOIS databases of the five Regional Internet Registries (RIRs: ARIN for North America, LACNIC for South America, RIPE NCC for Europe, AFRINIC for Africa, and APNIC for Asia/Pacific, including Australia) and the two National Internet Registries (NIRs: KRNIC for South Korea and JPNIC for Japan). This data contains the day that the data was collected, an organization ID, the last date that it was changed, an organization name, and a country.

The first challenge is that there are missing countries for many months. First, we group by organization ID and fill down then up within the organization. For example, in the below example the network is assigned C_a moving down to t_3 but C_b moving up to t_4 because the organization ID changed during that period. An alternative method could assume that the month is set until it definitively changes in the dataset.

$$\begin{bmatrix} AS\# & Month & OrgID & Country \\ AS_1 & t_1 & O_1 & C_a \\ AS_1 & t_2 & O_1 & ?? \\ AS_1 & t_3 & O_1 & ?? \\ AS_1 & t_4 & O_2 & ?? \\ AS_1 & t_5 & O_2 & ?? \\ AS_1 & t_6 & O_2 & C_b \end{bmatrix} \longrightarrow \begin{bmatrix} AS\# & Month & OrgID & Country \\ AS_1 & t_1 & O_1 & C_a \\ AS_1 & t_2 & O_1 & C_a \\ AS_1 & t_3 & O_1 & C_a \\ AS_1 & t_4 & O_2 & C_b \\ AS_1 & t_5 & O_2 & C_b \\ AS_1 & t_6 & O_2 & C_b \end{bmatrix}$$

After filling countries down and up, we then group by Autonomous System and fill countries down and up again. In the example below, the Autonomous System has no country data for any of the months where it was assigned to O_2, so after filling with the first method there would be two months with missing data. Each of those months for O_2 then take on the country value for O_1, the last time when we had definitive location data for the network.

$$\begin{bmatrix} AS\# & Month & OrgID & Country \\ AS_1 & t_1 & O_1 & C_a \\ AS_1 & t_2 & O_1 & ?? \\ AS_1 & t_3 & O_2 & ?? \\ AS_1 & t_4 & O_2 & ?? \\ AS_1 & t_5 & O_3 & ?? \\ AS_1 & t_6 & O_3 & C_b \end{bmatrix} \longrightarrow \begin{bmatrix} AS\# & Month & OrgID & Country \\ AS_1 & t_1 & O_1 & C_a \\ AS_1 & t_2 & O_1 & C_a \\ AS_1 & t_3 & O_2 & C_a \\ AS_1 & t_4 & O_2 & C_a \\ AS_1 & t_5 & O_3 & C_b \\ AS_1 & t_6 & O_3 & C_b \end{bmatrix}$$

Additionally, the CAIDA collects WHOIS data every three months, but the networks could change ownership or location any time during the months between the data dumps. This is a problem because interconnection (which is collected every month) could be erroneously assigned to one country when the registration had already changes. This could potentially occur for two months if the registration changed in the 30 days after the previous ASN to Organization dataset was published.

To correct against this, we purchased a subscription from Big Data Cloud (api.bigdatacloud.net), and ran it on all current ASN numbers to get registration details for all currently active systems. We then assign a new start and end date to the CAIDA dataset based on the individual registration data to get the precise date that it was registered. Many historical observations also contain information about when the registration was last changed - in these cases we alter the registration date to reflect the stated changed date. However, this information is not contained in all observations.

This results in a dataset with one observation for each Autonomous System for each country that the system was assigned to. Each observation has a start date, when the network was first assigned to the country, and an end date, when the country was last assigned to the country. There are a total of 105,926 unique networks in this dataset across 242 countries and territories. Not all of these networks have agreements during the entire period, and many were assigned during the past ten years. In the first month in the CAIDA dataset (January 1998) there were 3,221 networks with at least one connection, and in December 2020 there were 70,791 networks with at least one connection.

1.3 European Union Registered Networks

Internet service providers can also register as "European Union" located, which presents challenges for political scientists using interconnection data. We change the location of all networks listed as EU-located depending on the country where a majority of IP addresses assigned to the Autonomous System are located. We use data from "IPIP.net" to make this determination.[5] These 48 networks were re-assigned the countries in Table 1.

Table 1: Countries Gaining Autonomous Systems from "EU"

AT	BE	CH	DE	EG	ES	FI	FR	GR	IE	IS	IT	KW	MT
2	2	6	5	1	3	3	6	1	2	1	2	1	1
MZ	NL	NO	RO	SE	SN	UA							
2	21	1	1	1	1	1							

1.4 Sources of Potential Bias in Interconnection and Autonomous System Data

While differences in collection between the AS organization and AS relationshisp data is a potential source of measurement error, we have no reason to believe that this would be biased in any particular direction. It is possible that delays in assigning autonomous systems to countries might bias against interconnection to countries that are experiencing disproportionate increases in interconnection overall.

One significant challenge with this data is that networks may have a presence in multiple countries, which would not necessarily appear in this dataset. For instance, an company

[5] https://en.ipip.net/

based in the US may have an AS with endpoints outside of the US. This is a problem for only the largest systems, since the vast majority of ASes are geographically bounded to a single area. Groups headquartered in one country may also choose to register an AS in their home region rather than the region where the AS operates. We compare the registered country with the geographic spread of IP addresses, and change the location of autonomous systems which control no IP addresses in the country to which they are registered.

A second source of potential bias is that interconnection agreements are not perfect measures of data flows. These agreements imply that data can flow between two ASes, but not that data is flowing between the ASes. However, these agreements are highly correlated with self-reported traffic volume by ASes (Lodhi et al., 2014). However, these interconnection agreements do not all result in the same level of exchange. The following discussion of the data is in Zhuo, Huffaker, Claffy and Greenstein (2021).

> First, we note that we are able to capture only part of networks activities the formation and termination of interconnection agreements, and the types of agreements. It is important to note that connectivity is not traffic, though there is evidence that IP address space advertised by BGP tables are strongly positively correlated with networks self-reported traffic volume for a large set of peer-to-peer interconnections (Lodhi et al., 2014). We do not know how much traffic exchange happens across an interconnection or how that traffc has changed over time. If major changes in traffic occurred purely through existing interconnections, causing increased or decreased investment in Internet infrastructure, it would be invisible in our data.

2 Independent Variables

2.1 Alliances

Figure 1 contains the country-treaty observations where the status of the treaty changed between January 2010 and December 2017. It does not include all alliances - only treaties which were either entered into or withdrawn from during the study period. This data is taken from the *Alliance Treaty Obligations and Provisions Project (ATOP)* data on military agreements (Leeds, Ritter, Mitchell and Long, 2002). Version 5.0 of the dataset covers all alliances formed between 1815 and December 31, 2018.

Figure 1: Treaties Changing Status

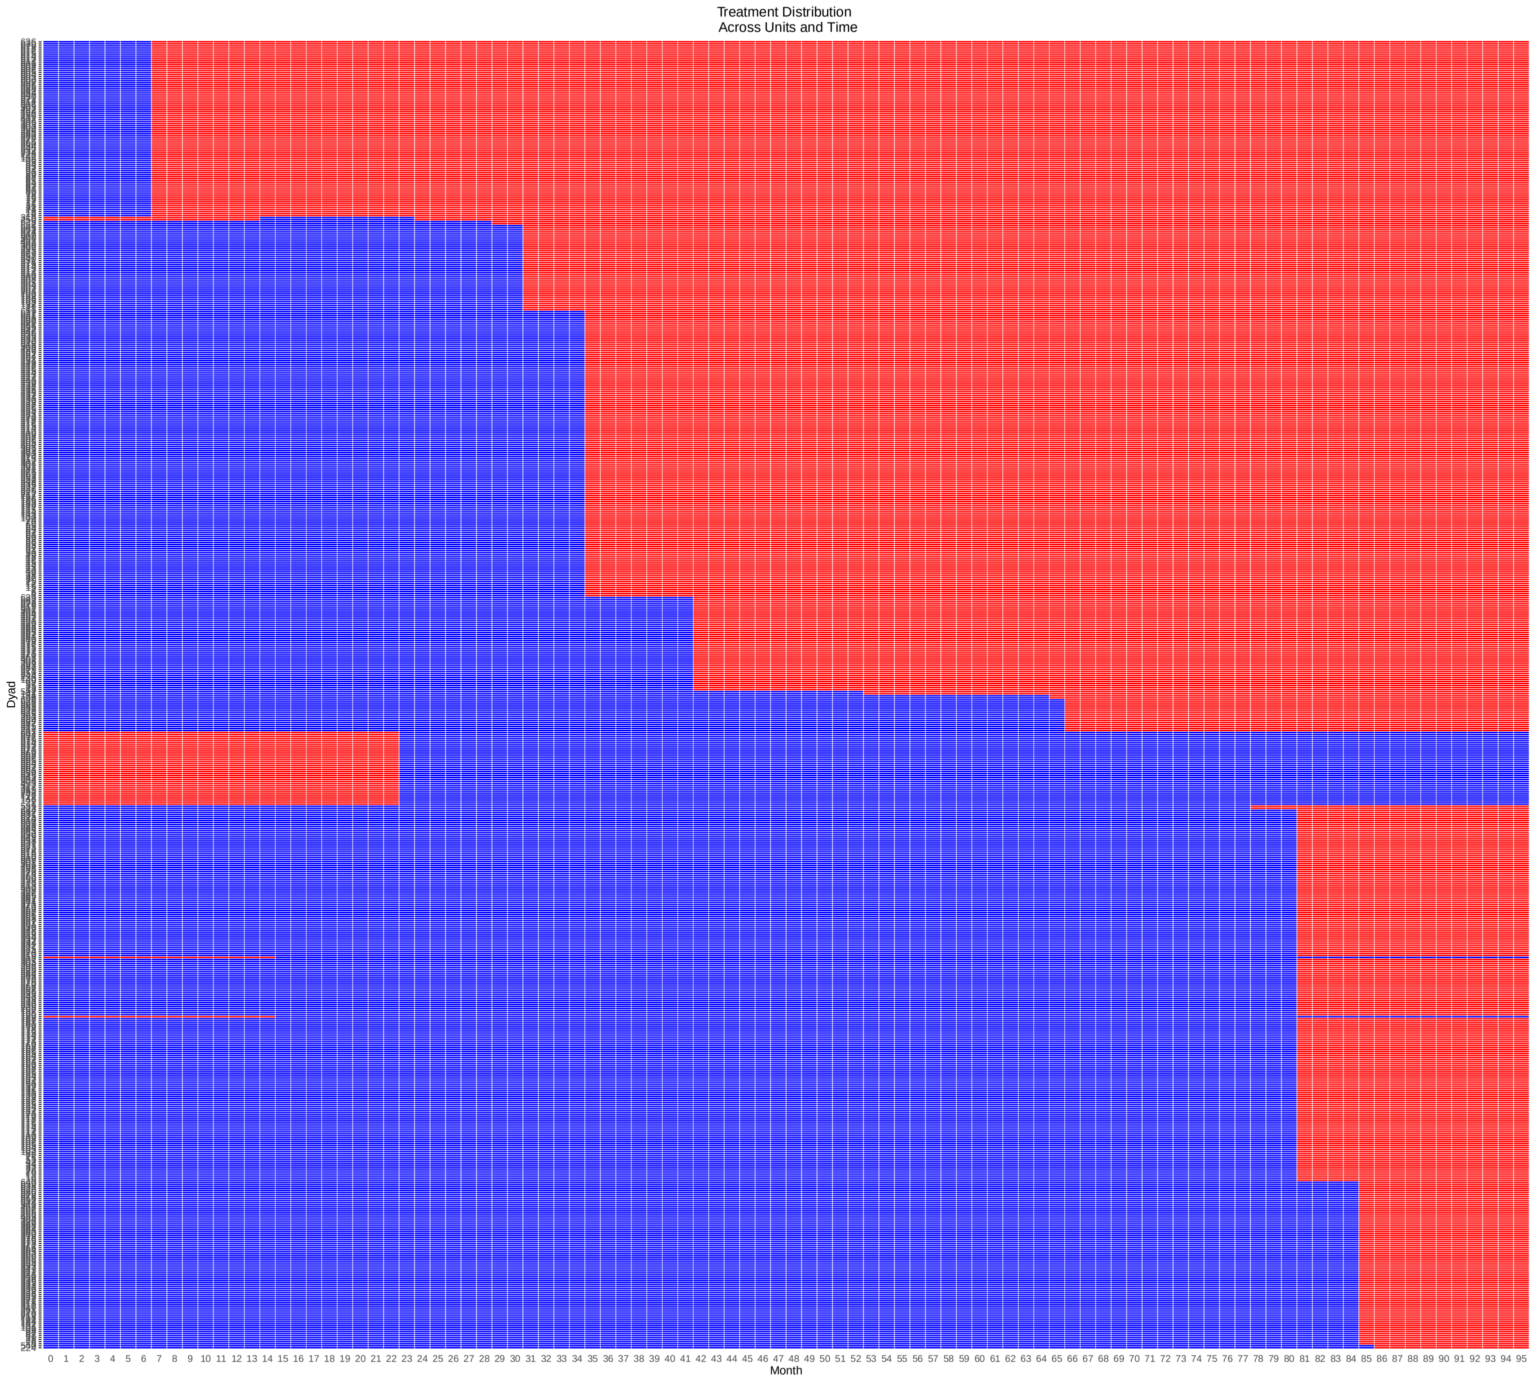

Figure 2: Proportion of Dyads Covered by Alliance or Treaty

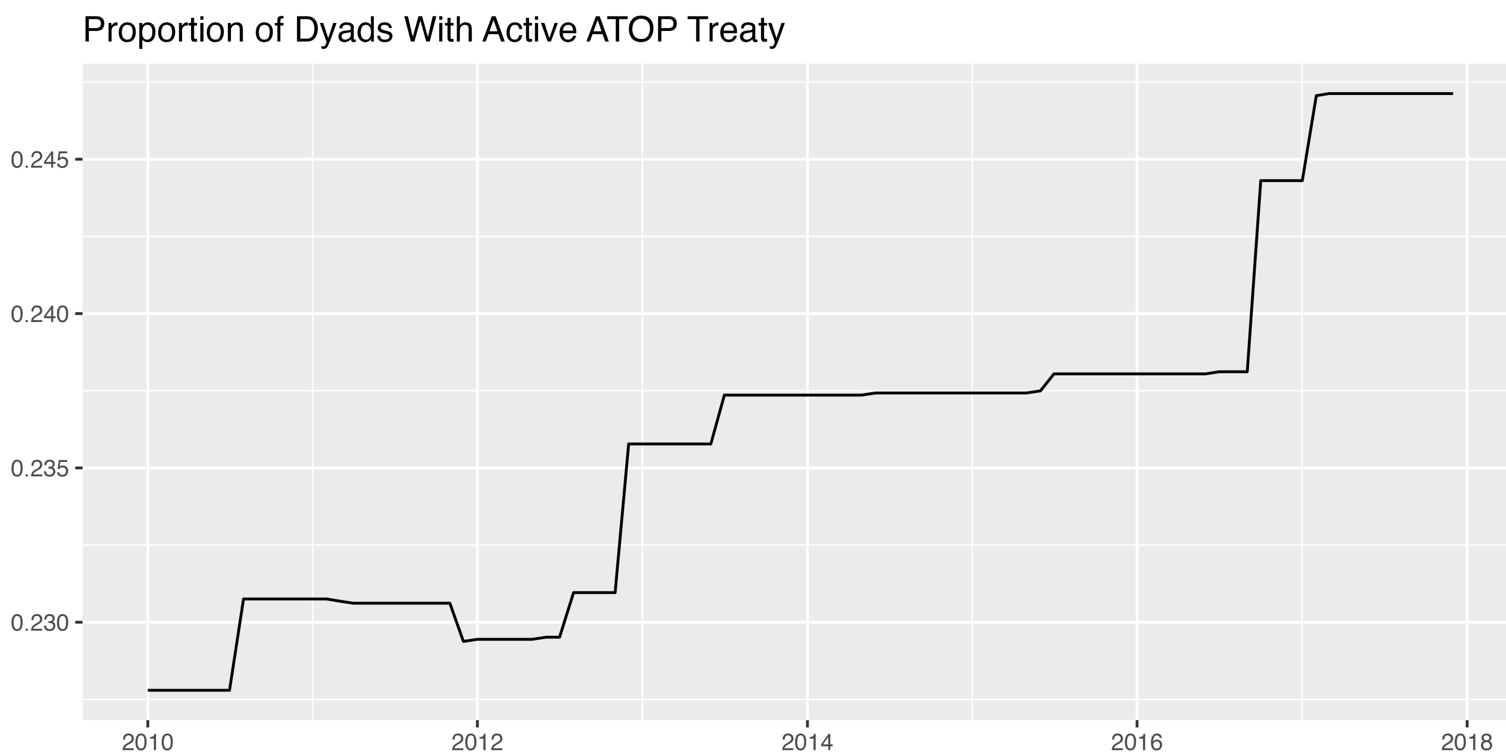

2.2 Control Data

I control for two factors that varied significantly at the dyad-level during the study period. The first is joint WTO membership - eleven countries joined the WTO during the study period. All dyads featuring one of the countries in the table below varied in the joint WTO membership control.

Several states entered into bilateral or multilateral preferential trade agreements during this period - I control for this using the *Design of Trade Agreements (DESTA)* dataset Version 2.0 for the preferential trade agreements (Dür, Baccini and Elsig, 2014).

Table 2: WTO Entries 2010-2018

Country	Date of Entry
Afghanistan	29 July 2016
Liberia	14 July 2016
Kazakhstan	30 November 2015
Seychelles	26 April 2015
Yemen	26 June 2014
Tajikistan	2 March 2013
Laos	2 February 2013
Vanuatu	24 August 2012
Russia	22 August 2012
Samoa	10 May 2012
Montenegro	29 April 2012

3 Fixed Costs of Interconnection

One significant fixed costs for internet exchange is fiber-optic cables that carry data across long distances. To account for this potential bias we first collect data on submarine and terrestrial cable networks. For data on current submarine cables I use information from Telegeography,[6] and for data on unused, or "dark" cables I use data from the Submarine Cable Almanac, which began in 2011.[7] Each cable contains information about the set of physical landing points where the cable connects to terrestrial networks, which I project into a adjacency matrix of countries.

However, multiple cables can meet in one landing point, which gives an opportunity to easily exchange data between the systems. This is another way of accounting for the fixed data exchange costs. For example, one country might pursue a cable with another to a landing point where data can then flow through the other cables that meet at that point. I create an adjacency matrix for cables with shared landing points, and use this to create another adjacency matrix of countries that can exchange data through one direct submarine path through two cables.

Furthermore, terrestrial cable networks allow countries to exchange data over ground. For this reason, all contiguous countries are also selected, since terrestrial fiber networks are not mapped and available the same way as submarine cables. I assume that if two countries are contiguous they do not have significant fixed costs to move data.[8] Finally, there are terrestrial networks that connect multiple countries. This is particularly important for European landlocked countries such as Switzerland and Austria, which are highly integrated into the global internet but do not have submarine cables. Russia also relies heavily on terrestrial cable networks to Europe and Asia.

[6] TeleGeography. "Submarine Cable Map." Submarine Cable Map. https://www.submarinecablemap.com/.

[7] Submarine Telcoms Forum. "Submarine Cable Almanac." https://subtelforum.com/products/submarine-cable-almanac/.

[8] For example, see the International Telecommunications Union map of voluntarily disclosed terrestrial networks. https://www.itu.int/itu-d/tnd-map-public/

I include three terrestrial fiber networks in the analysis, TTK Eurasia Highway, the TEA Cables, and the European fiber network. European networks include a variety of interconnected internet backbones including the Pan-European Crossing.[9] Countries in Europe have been highly interconnected since before the study period in 2008 (Rutherford, Gillespie and Richardson, 2004). TTK Eurasian highway has connected Europe and Asia through Russia since before the study period as well.[10] Additionally, Chinese and Russian operators have invested in the TEA Cable network to move data since 2010.[11]

Table 3: Terrestrial Cable Networks

Cable Network	Countries
TEA Cable	SE, FI, RU, CN, JP, HK, DE, UA, FR, NL, SE
TTK Eurasia	CN, MN, JP, FI, EE, LT, LV, PL, RU, NL
European Backbone	BE, BG, CZ, DK, DE, EE, IE, ES, FR, HR, IT, CY, LV, LT, LU, HU, MT, NL, AT, PL, PT, RO, SI, SK, FI, SE, BY, CH, GB, NO

Figure 3 contains the network of dyads connected by existing fiber networks during the first month of the study.

Second, we collect information on internet exchange points (IXPs). This data comes from Packet Clearing House (PCH) through its website database.[12] We individually validated both the IXPs functionality and its founding data. We identified IXPs which were either never operational or shut down before 2009. In total, we exclude sixteen IXPs, contained in Table 4. Some IXPs have no record of existing outside PCH (such as Cambodia's Finder INternet Exchange), while others existed but we were unable to identify whether they remained operational by 2010 (such as Eswatini's Swaziland Internet Exchange). This process only changes the sample when the IXP is the first for a country, since we select dyads based on whether they have at least one IXP.

[9] "Global Crossing Expands Pan European Network; Secures Additional Rights of Way," March 9, 1999. https://www.fiberopticsonline.com/doc/global-crossing-expands-pan-european-network-0001.

[10] Totaltelecom. "TTK Triples International Data Transit Capacity between Europe and Asia," July 27, 2009. https://www.totaltele.com/447602/TTK-triples-international-data-transit-capacity-between-Europe-and-Asia.; Submarine Cable Networks. "ERA and HSCS Broaden the Eurasia Highway," July 15, 2011. https://www.submarinenetworks.com/systems/asia-europe-africa/hscs/era-hscs-eurasia-highway.

[11] "Rostelecom: Transit Europe-Asia The New Opportunities." Moscow, October 27, 2011. https://www.hkcolo.com/hkconnect/2011/event/thankyou/pdf/Rostelecom_Irina.pdf.

[12] https://www.pch.net/

Figure 3: Network of Dyads With Submarine Cable Interdependence

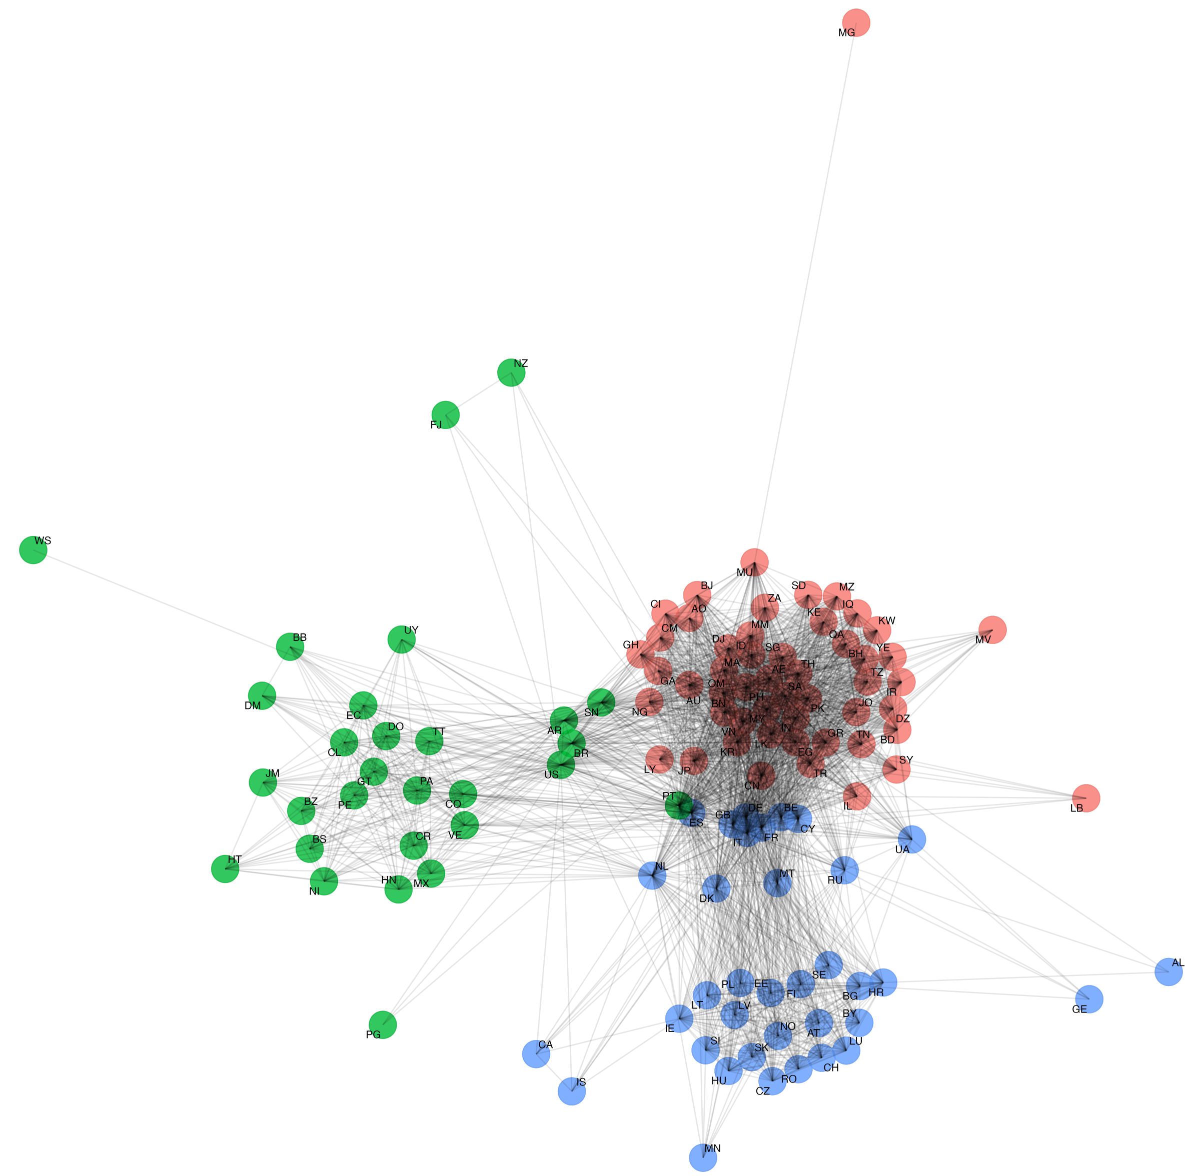

Table 4: IXPs Excluded from Analysis

Name	PCH Status
Nicaraguan Internet Exchange	Unknown
NAP Perú	Unknown
Swaziland Internet Exchange	Deprecated
Bangladesh Society of Internet Exchange	Deprecated
Internet Exchange of Saudi Arabia	Unknown
PKIXP Karachi (ex ZPIX)	Unknown
Kyrgyz IXP	Unknown
Côte d'Ivoire Internet Exchange Point	Defunct
Bucharest Internet Exchange	Defunct
IX de Bolivia	Defunct
BurundiX Internet Exchange Point	Defunct
CAPADI NAP-PY	Defunct
Zimbabwe Internet Exchange	Defunct
Tajik Internet Exchange Point	Defunct
Common Routing Exchange	Defunct
Finder Internet Exchange	Defunct

4 State-Owned Enterprises

Data on whther networks are operated by state-owned enterprises comes from Carisimo et al. (2021). Here I present more information on these networks and their distribution in the data. The table below contains the countries with four or more state-owned networks. Unsurprisingly, the countries with the most state-owned are China and Russia. I lookup basic information on these networks from BigData Cloud (`https://www.bigdatacloud.com/`). The networks with the most interconnection agreements with significant state-ownership are operated by Angola Telecom, Telekomunikasi Indonesia, Emirates Telecommunications Group, Swisscom, and China Mobile. The networks with the most addresses are operated by Chinanet, China Unicom, China Mobile, Rostelecom (Russia), and Telecom Egypt.

Table 5: Countries with Four or More State-Owned Networks

China	228
Russian Federation	181
Singapore	88
Norway	39
United Arab Emirates	33
Qatar	20
Vietnam	17
Saudi Arabia	15
Thailand	14
Malaysia	13
Serbia	13
Belgium	12
Tunisia	12
Bangladesh	11
Iran (Islamic Republic of)	9
South Africa	9
Uzbekistan	9
Fiji	8
Indonesia	8
Kazakhstan	8
Angola	7
Ecuador	7
Switzerland	7
Azerbaijan	6
Bahrain	6
Belarus	6
Libya	6
Papua New Guinea	6
Colombia	5
Hungary	5
Micronesia (Federated States of)	5
Mozambique	5
Slovenia	5
Algeria	4
Cuba	4
Pakistan	4
Sri Lanka	4
Zimbabwe	4

5 Results: Linear Model

Table 6 models interconnection as a linear form with $\log(y+1)$ as the dependent variable.

Table 6: Results with Linear Model

Dependent Variable:	**Agreements (log+1)**			
Alliance	0.0937***	0.0937***	0.0935***	0.0935***
	(0.0231)	(0.0231)	(0.0231)	(0.0231)
Joint WTO		0.0333***		0.0333***
		(0.0068)		(0.0068)
PTA			-0.0148	-0.0148
			(0.0096)	(0.0096)
Observations	2,790,720	2,790,720	2,790,720	2,790,720
Fixed-effects				
CountryA*Month	Yes	Yes	Yes	Yes
CountryB*Month	Yes	Yes	Yes	Yes
Dyad	Yes	Yes	Yes	Yes

Clustered (Dyad & Month) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

6 Results: Differenced Model

Table 7 models interconnection as a linear form with the dependent variable differenced at the dyad level. The fixed effect in this case is CountryA*Month and CountryB*Month, dropping the dyad fixed effect in the previous results.

Table 7: Results with Differenced Outcome Variable

Dependent Variable:	$Agreements_t$ - $Agreements_{t-1}$			
Model:	(1)	(2)	(3)	(4)
Variables				
Treaty	0.1127***	0.1127***	0.1425***	0.1425***
	(0.0294)	(0.0294)	(0.0412)	(0.0412)
Joint WTO		-0.0259		-0.0258
		(0.0395)		(0.0395)
PTA			-0.0616**	-0.0616**
			(0.0288)	(0.0288)
Fixed-effects				
CountryA*Month	Yes	Yes	Yes	Yes
CountryB*Month	Yes	Yes	Yes	Yes

Clustered (paird & time_seq) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

7 Robustness Checks

This section presents a robustness check for the main findings in the paper. I check for robustness by sequentially dropping individual dyads from the analysis to understand when one country has the potential to change the results in the paper. The figures present the point estimate for the main effect along with a 95% confidence interval. None of the results meaningfully changed by removing treated dyads sequentially.

Figure 4: Effect of Alliances Dropping Treated Dyads With Replacement

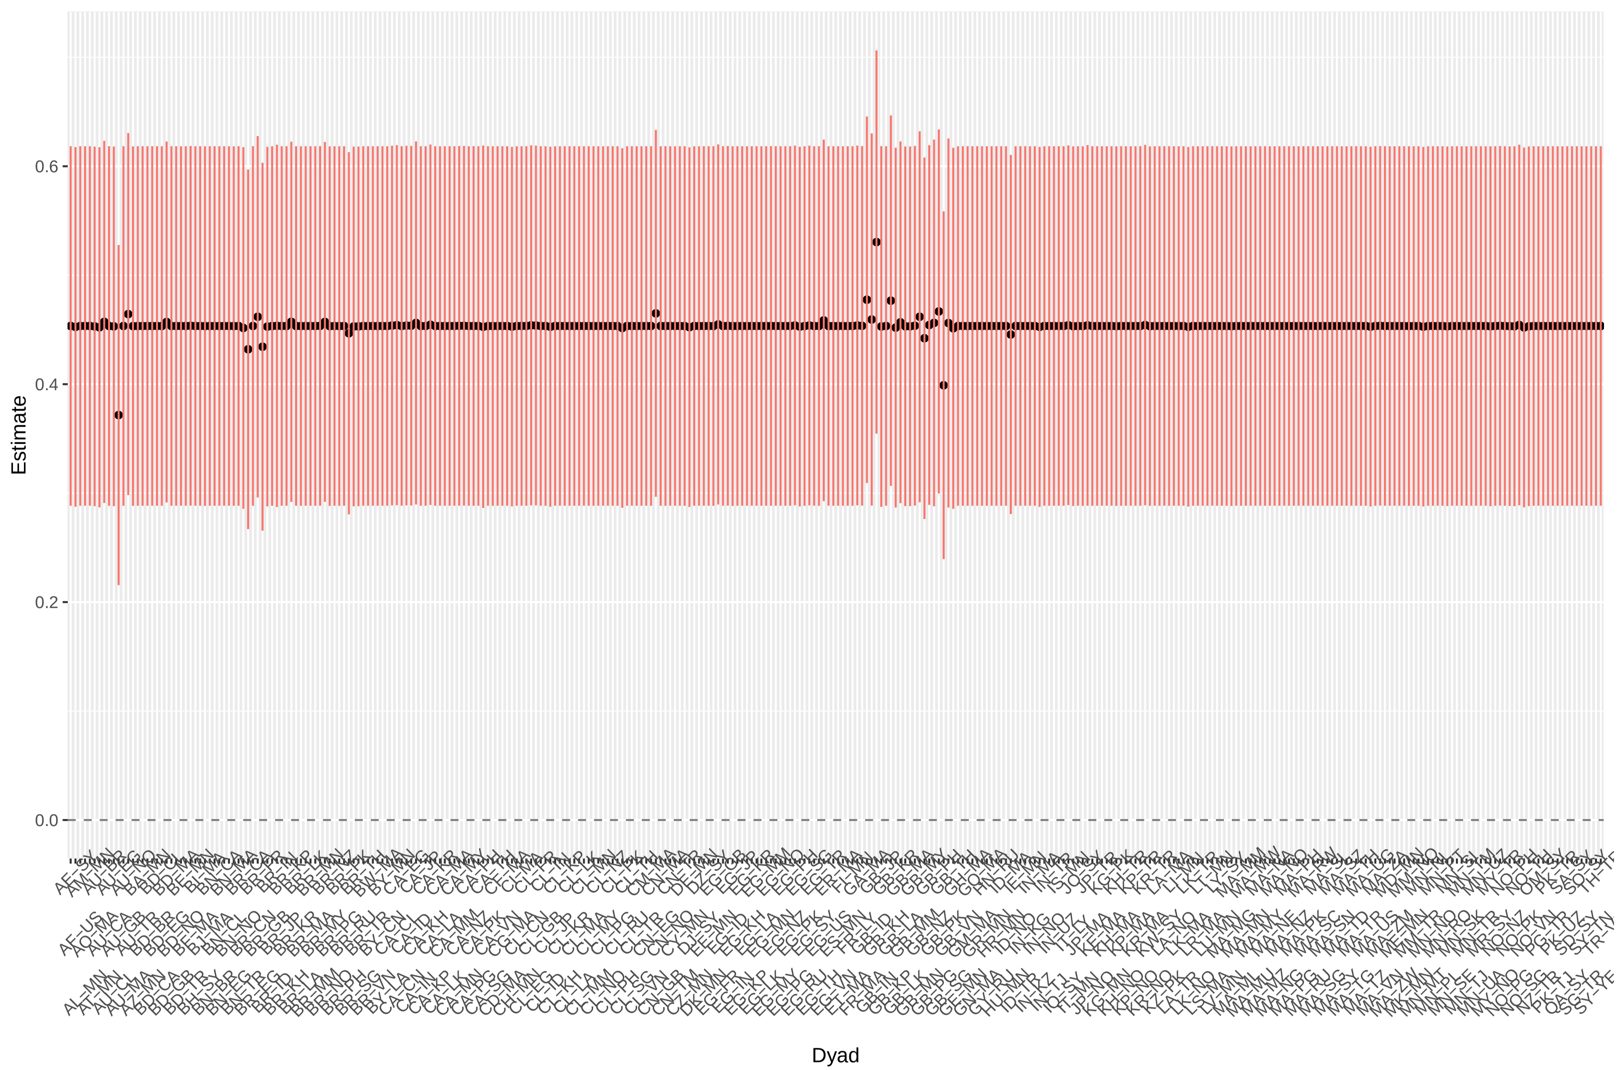

This figure presents the results from the main paper dropping each alliance dyad (with replacement).

8 Sanctions Discussion in SEC Filings

Cable & Wireless is a subsidiary of Liberty Latin America Ltd., which is headquartered in Denver, CO and incorporated in Bermuda. The main ASN they operate has a majority of IP addresses located in the United States, and they have customers in both Cuba and Venezuela. Their SEC filing 10k discusses their compliance with US sanctions and OFAC risks.

2023 10K Form (https://www.sec.gov/Archives/edgar/data/1712184/000171218424000030/lila-20231231.htm)

> For example, certain of our companies provide (and may in the future provide), directly or indirectly, certain services to governmental entities in Cuba (e.g., C&W sells IP and international transport telecommunication services to ETECSA, the Cuba state-owned telecommunications provider and to other international telecommunications providers that in turn sell telecom services to ETECSA). All these services are provided outside of Cuba and the provision of non-facilities based telecom services to Cuba are permissible under the Cuba Assets Control Regulations and a general license from OFAC.
>
> We also have interconnection and services contracts with telecommunications carriers located in Venezuela. With respect to Venezuela, we have advised OFAC that we believe that our activities there are not covered by the OFAC regulations or are otherwise allowed under a general license and exemptions or, in the alternative, should be licensed by OFAC. In September 2022, OFAC issued a specific license to allow us to engage in all transactions necessary for U.S. financial institutions to process the collection of outstanding debts and the receipt of current and future payments relating to telecommunications services provided to Compañía Anónima Nacional Teléfonos de Venezuela. OFAC extended this license on August 17, 2023.

References

Carisimo, Esteban, Alexander Gamero-Garrido, Alex C. Snoeren and Alberto Dainotti. 2021. Identifying ASes of State-Owned Internet Operators. In *Proceedings of the 21st ACM Internet Measurement Conference*. Virtual Event: ACM pp. 687–702.

Dür, Andreas, Leonardo Baccini and Manfred Elsig. 2014. "The Design of International Trade Agreements: Introducing a New Dataset." *The Review of International Organizations* 9(3):353–375.

Huston, Geoff. 2016. "The Death of Transit?".

Kolkman, Olaf, Andrei Robachevsky, Carl Gahnberg and Hosein Badran. 2022. Evolution of the Edge, What about the Internet? In *Proceedings of the ACM SIGCOMM Workshop on Future of Internet Routing & Addressing*. Amsterdam Netherlands: ACM pp. 1–5.

Leeds, Brett Ashley, Jeffrey M. Ritter, Sara McLaughlin Mitchell and Andrew G. Long. 2002. “Alliance Treaty Obligations and Provisions, 1815-1944.” *International Interactions* 28:237–260.

Lodhi, Aemen, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis and KC Claffy. 2014. “Using peeringDB to Understand the Peering Ecosystem.” *ACM SIGCOMM Computer Communication Review* 44(2):20–27.

Rutherford, Jonathan, Andrew Gillespie and Ranald Richardson. 2004. “The Territoriality of Pan-European Telecommunications Backbone Networks.” *Journal of Urban Technology* 11(3):1–34.

Zhuo, Ran, Bradley Huffaker, KC Claffy and Shane Greenstein. 2021. “The Impact of the General Data Protection Regulation on Internet Interconnection.” *Telecommunications Policy* 45(2).