International Trade in Data: Investment and Usage of the Subsea Internet Cable Network*

Jihye Jeon† and Marc Rysman‡

January 2024

Abstract

This paper studies the construction and use of undersea internet cables, a critical piece of communication infrastructure. We view traffic on these cables as *international trade in data*. We propose a model in which country-to-country trade in data, similar to a gravity equation model, traverses the cable network leading to interregional flows observed in our data. On the supply side, firms decide whether to invest in new cables, recognizing the impact of their investment in any one market on global internet flows. We estimate this model using new data on cable construction and usage via moment inequalities. We use the results to decompose growth in global internet usage into growth in demand and improvements in the cable network. We find that the latter is an important contributor on par with the former. Our counterfactuals highlight the role of business stealing and network externalities in generating inefficient allocation of cables.

*We thank Haoran Pan and Tim Lee for excellent research assistance. We also thank Savannah Kelly, Lilian Lea, and Paul Brodsky from Telegeography for providing valuable insights about the industry and the data.

†Boston University. Email: jjeon@bu.edu

‡Boston University. Email: mrysman@bu.edu.

1 Introduction

Communication technology is critical to modern trade and economic activity. A central piece of communication infrastructure is the network of subsea fiber optic cables that carry internet data. More than 99% of international data traffic travels through these subsea cables, which remain the most efficient way to send information across the ocean. Clark (2018, pg. 24) writes about fiber optics: "The economic implications of its development and deployment are perhaps the single most important factor in the growth and success of the internet." Growth has indeed been impressive; global internet traffic grew by a factor of 1000 between 2002 and 2020, far surpassing the growth of traditional trade of goods and services.

We study the development, growth, and usage of the subsea internet cable network to shed light on this understudied yet important market. We leverage new data on the subsea cable network to study two main research questions. First, we characterize flows on the network as *international trade in data*. Studying trade in data provides a new perspective on globalization and the sources of economic growth. We develop an analog of a gravity model for trade in data that accounts for both country-to-country demand to exchange data and the quality of the network that connects countries. We decompose the relative contribution of these two factors in the evolution of trade in data.

Second, we study investments to build new cables. A new cable in one region can affect traffic flows globally. The networked nature of this product means that investors' private incentives can deviate substantially from the social optimum. In addition, the subsea internet cable market is subject to a very low level of economic regulation. As a result, the network we observe is unlikely to exhibit global efficiency. We develop a model that highlights the international implications of local investments and the role of inefficient business stealing incentives in investment decisions. Furthermore, whereas cables have traditionally been owned by telecommunication carriers such as AT&T and British Telecom, recent construction has been dominated by content providers such as Google and Microsoft. We allow parameters in the objective function of owners to differ between telecoms and content providers to study how their incentives differ.

Our data is primarily drawn from Telegeography, a proprietary data company that covers internet equipment and services, with particular expertise in the sub-

sea cable industry. We observe landing points and construction dates of all cables, and typically more information such as bandwidth capacity, length of the cable, and owners and ownership shares. Critically, we observe used bandwidth, which is a measure of usage of a cable that we further describe below. We observe used bandwidth between some regions and also at the level of the individual countries. Our current data set goes from 2002 to 2020. We supplement the data with various sources, such as data on internet penetration from the International Telecommunication Union and trade data from CEPII.

We use our data to document several patterns in cable investment and usage. First, data flows on these cables have grown at a dramatic pace, and they are heavily concentrated geographically. Second, we show the slow arrival of cables in certain regions in contrast to repeated investments in enormous quantities in others. Third, we show new cable entry is associated with market expansion but its effect diminishes with the number of existing cables or existing capacity, suggesting that the geographical imbalance in cable construction may have important welfare implications. Lastly, we document the growth in the share of capacity invested by content providers. The investment patterns we observe in the data may be simply due to heterogeneity in demand, but they may also be explained by other factors such as business stealing incentives and externalities across markets. We propose a model that allows us to distinguish these factors.

We face an empirical challenge in that we observe a measure of how much data flows between two regions, but we do not observe endpoint-to-endpoint usage. That is, data flowing between two countries may have different starting and endpoint countries, and data traveling between two endpoint countries often take multiple routes as determined by various factors, such as cable availability, capacity, and length. To address this issue, we develop a structural model that maps country-to-country demand for data into observed data flows. We model bilateral data demand through a gravity equation in which demand is a function of characteristics such as the GDP and population of the two countries. The cable network determines the set of available paths that data may take between two countries and the features of those paths, which determines the quality of the connection. We model the allocation of demand for data across these different potential paths. In this way, the existing cable network determines which cable paths data travel on and thus, the observed usage of the cables.

On the supply side, we develop a model of cable construction. In each period, opportunities to build new cables in markets (i.e. pairs of regions) appear exogenously, and investors choose whether to join consortiums that can take advantage of these opportunities. An investor's decision-making in one market depends on demand and its own and its rivals' cable ownership states not only for that particular market but also for all other markets. Firms profit based on their ownership share of cables and the amount of data that traverses those cables. Our model highlights the business-stealing nature of cable construction. In our model, new cables on busy routes, which typically have substantial previous investment, often have high private value relative to social value as new owners steal business from existing cables. In contrast, cables on less busy routes, which typically have lower investment levels, are the reverse: the cable may have high social value but relatively low demand.

We estimate our demand model, drawing moments from the variation in usage across continents and years. We also exploit "island countries," for which all communication must traverse subsea cables. The estimates suggest that a one-percent increase in the sum of the GDP of two countries involved in trade would result in a four-percent increase in trade in data. Further, the path length has a negative effect, and the path capacity and the number of cables have positive effects on the amount of data flows allocated to the path.

We use the demand model to pin down quantities around the world with and without investment. Combined with construction costs, which we observe in our data, we estimate markups (that is, variable profit) utilizing moment inequalities. We recover a 95% confidence interval for the variable profit for four distinct time periods separately and for content providers and non-content providers (including carriers and other private investors) separately. For example, we estimate that a one-Tbps increase in bandwidth served by the firm would result in a profit increase between 11 to 12 million dollars for non-content providers in 2018-2020.

We find that the markups fell rapidly from 2002 to 2020, consistent with the steep decline for bandwidth pricing also observed in this market. We also find that the markups are much higher for non-content providers. The gap was the largest in 2008-2012 when Google built its own first cable, but it narrowed significantly over time. This is consistent with the fact that content providers have different incentives for building cables, such as connecting their own data centers, and that they are not

allowed to sell capacity directly.

Taking the estimates together, we conduct two main counterfactuals. First, we decompose changes in data flows into changes in demographics and changes in the quality of the cable network. We find that cable construction was a substantial contributor to the growth of trade in demand, on par with increases in demand. In 2014, for example, improving the quality of the cable network to the 2020 level while fixing the demographics at the observed level would double the total trade in demand from the level observed in 2014, or bring it up to 36% of the 2020 level.

Our second set of counterfactuals considers a de novo entrant adding a new cable in each market to evaluate the efficiency in the cable network. We compute the change in the combined profits of incumbent firms, decomposing it into the change in the market where the investment happens and the change in the rest of the world. It is ambiguous how a new cable construction would affect incumbent profits. The new cable construction would shift market shares from incumbents to the entrant in the market with the investment, decreasing the incumbent profit in that market. At the same time, the increased capacity would attract more data flows into cables in that market, shifting down the incumbent profit in the other markets. Lastly, the increased capacity would have a market expansion effect, typically in multiple markets, with increased demand served by cables across different markets. This would increase incumbent profits inside and outside the investment market.

Building a new cable has a negative effect on incumbent profits on average. However, the effect varies widely across markets. For example, we find that the effect is large and negative in markets that involve North America, which already possess large capacity. This suggests that substantial business stealing happens in these markets. By contrast, incumbents would experience profit increases in many of the markets involving Sub-Saharan Africa, suggesting that a new cable would have substantial market expansion effects in these markets. Meanwhile, entrants would enjoy higher profits from investing in the North American markets, which indicates that their private incentives may be misaligned with socially optimal outcomes.

Related Literature

There is little economic research on the subsea cables market. Hjort and Poulsen (2019) show that the arrival of subsea cables to Africa improved internet service for

consumers living along the relevant terrestrial network and led to increased productivity and high-skilled employment. Their paper highlights the economic impact of internet access driven by the expansion of the subsea cable network, but does not explore when or why cables are constructed, whereas our work focuses on the market for subsea cables itself. The closest paper that we are aware of is the concurrent paper by Caoui and Steck (2023), which focuses on the role of path diversity. The main difference from their paper is that our model takes into account the fact that all regions are globally connected, which means that demand for communication in one market can be transmitted through cables in other markets, and new construction in one market can affect data flows in others. This is a key feature not only in this setting but also in other settings such as transportation markets (e.g., shipping ports, airports, and railroads), power and energy industries, and other telecommunication markets (e.g., data centers). More broadly, our paper relates to papers that study infrastructure that supports the internet such as Greenstein (2015).

Our paper also connects to papers on the interconnection between the internet and trade. Steinwender (2018) studies how advances in telecommunications technology affect trade. Blum and Goldfarb (2006) examines whether the law of gravity holds in the case of digital goods sold over the internet.

Our approach is related to equilibrium models of transportation and cargo networks, such as Fréchette et al. (2019) Brancaccio et al. (2020). Allen and Arkolakis (2022) study driving traffic on a road network which has some similarity to our approach. Our research is also related to empirical research on the construction of networks in retail, such as Jia (2008), Holmes (2011), and Houde et al. (2022), and in airlines, such as Aguirregabiria and Ho (2010) and Li et al. (2022). Our research also relates to papers on investment in trade-related infrastructure such as Jeon (2022) and Brancaccio et al. (2020) and in internet services markets such as Rysman (2016). A related paper on investment in telecommunication infrastructure is Elliott et al. (2023).

Lastly, we contribute to a growing empirical literature that employs moment inequalities, recently reviewed by Kline et al. (2021) and Canay et al. (2023). Our paper is especially closely related to papers that estimate an endogenous entry or product choice model using inequalities such as Eizenberg (2014), Wollmann (2018), Dickstein and Morales (2018), and Fan and Yang (2022).

The rest of the paper proceeds as follows. Section 2 provides background on the

subsea cables market, information about our data, and motivating facts. Section 3 describes the model of demand for international data, cable usage, and cable construction. Section 4 discusses our estimation strategy and results. Section 5 presents results from counterfactual exercises. Section 6 concludes.

2 Subsea Cable Networks and Data

2.1 Subsea Cable Networks

As of 2021, there are over 400 active subsea cables spanning 1.5 million kilometers in length, connecting 182 countries across all continents except Antarctica. Industry experts estimate that the undersea telecommunication cable network carries about 95% of intercontinental global internet traffic, and 99% of transoceanic digital communications such as voice and data including trillions in international financial transactions daily (CRS, 2022). The core of a cable consists of strands of glass fibers through which lasers can propel data, which is surrounded by additional layers that are meant to provide protection and transfer power. A cable is buried on the ocean floor by specialized cable-laying ships through a long and complex process.

Typically, cables are owned by a consortium of owners, including telecommunication companies, governments, and other businesses. Historically, telecommunication providers such as AT&T, British Telecom, and NTT (of Japan) were the main owners of the cables. Starting around 2017, an overwhelming share of investment has been by tech firms or content providers such as Google and Microsoft that are seeking to meet the growing demand for data exchanges and more control over ownership and usage.

Cable owners sell capacity on their lines to local carriers. Within the United States, a standard supply chain would feature cable owners selling capacity to internet transit providers such as Lumen Technologies (formerly Level 3), who would then sell internet service to local access providers such as Verizon or Comcast. In practice, there is some vertical integration between participants at each of these levels. In some foreign countries, internet service at both the local access and transit levels are provided by a national telecommunication firm, which may also participate in subsea cable construction.

The lack of economic regulation is a striking feature of the undersea internet cable market. There is some government intervention in regard to national security, but that has little impact on the features of the market we consider. For example, in the United States, all subsea cable landings must be approved by the Federal Communication Commission (FCC). Historically, this power was used to ensure reciprocal acceptance of cables owned by US entities in foreign countries. However, the FCC has not used its powers for this purpose in several decades. Rather, the FCC refers applications to Team Telecom, an ad hoc committee of interested federal agencies.[1] The Department of Homeland Security often takes the lead in these instances, particularly monitoring landings for their potential for terrorist attacks, as well as the involvement of Chinese owners and equipment makers. We regard this form of regulation as non-economic. While these national security considerations may affect the way a landing point is constructed, the exact location of a landing point within a broader region, and the choice of equipment, these are not features we study. Thus, we model the market as unregulated.[2]

2.2 Data

The main data we use in this paper are from Telegeography, a telecommunications market research company that provides detailed information on subsea cables and international bandwidth usage. Telegeography provides a comprehensive picture of the subsea cable network based on information collected using various methods, including confidential surveys, interviews with telecommunications company executives and engineering staff, and other network discovery tools such as aerial or satellite photographs. We are not aware of previous research in economics making use of these data, other than the concurrent paper by Caoui and Steck (2023). Telegeography's data provide the characteristics of active and planned subsea cables, including ready-for-service year, cable length, construction cost, ownership structure, landing points, and various capacity measures. We do not observe ter-

[1]The Trump Administration turned Team Telecom into a formal committee called the Committee for the Assessment of Foreign Participation in the United States Telecommunications Services Sector in 2020. Mergers in which both merging firms have holdings of subsea internet cables, such as the merger of Level 3 and Global Crossing, involve economic analysis of this market by the FCC. However, we are not aware of merging parties facing issues with their subsea holdings.

[2]Much of the information in this paragraph is drawn from an interview of a knowledgeable economist at the FCC.

restrial cables. Figure 1 provides a map from Telegeography of commercial subsea cables in 2023.

We observe several measures of capacity. *Potential capacity* is the capacity of data that a cable could potentially carry. Activating that capacity is costly, and *lit capacity* is a measure of how much of potential capacity is ready for service. Lit capacity is what is available to cable owners to sell. Primary purchasers are local carriers and content providers, although some large enterprises (e.g., governments, large companies) and educational institutions (e.g. CERN) also take a share. *Purchased capacity* is the amount sold to customers. After purchasing, customers choose how much to activate, which requires further investment. We observe this as well, as *used capacity*. Telegeography constructs an equivalent to used capacity for content providers that own and use their own cable. Note that whereas a retail consumer may purchase a certain amount of data communication per month, cable owners and their direct customers always transact in terms of capacity, for example, five wavelengths on a fiber optic cable for five years.[3] Traffic is the amount of data actually transmitted over the cable, but we do not have measures of this. We take used capacity as our measure of usage. In general, purchasers of capacity devote substantial engineering resources to tightly manage the relationship between purchased capacity, used capacity, traffic, and realized quality of communication, so used capacity is likely a good approximation of traffic.

Telegeography offers bandwidth price data at the city-pair-level. We do not use this data in our analysis for several reasons. The pricing data cover only a small subset of city-to-city pairs (even country-to-country pairs), which would limit our ability to study the global market. These city-to-city prices emerge from a market downstream from cables and typically cannot be linked to the use of any particular cable or cables, and it is unclear how much quantity that traverses cables actually pays these prices. For instance, content providers would not pay these prices for traffic on their own cables even though their traffic is included in our usage measure.

It is often important to distinguish between the number of cables in an area and the capacity of those cables. For example, Indonesia is a country of islands with many subsea cables connecting many relatively small islands and is a world leader

[3] A strand of fiber optic wire may carry 100 wavelengths of light. A fiber optic cable typically carries from 4 to 24 wires during our sample period.

in terms of the number of cables landing on its shores but is only 29th in the world in terms of total capacity of its landings in 2020.

Telegeography collects data on subsea cables, not terrestrial cables, and we are not aware of systematic data on terrestrial cables. As a result, we focus our analysis on communication between countries that primarily use subsea routes. We define *regions* such that communication between countries in different regions must traverse subsea cables for at least part of the way. Countries in the same region typically have a land route between them and a large share of communication goes over terrestrial cables only.[4]

We divide the world into the following seven regions: East Asia, Europe, North America, Oceania, South America, South Asia, and Sub-Saharan Africa. Telegegraphy has a finer definition of regions (13 regions) in their original data. After discussions with Telegeography, we combine certain regions into one. For example, we group Central Asia & Caucasus, Eastern Europe, and Western Europe into one region of Europe based on the observation that there are many viable terrestrial connections across the four regions. By contrast, there is almost no terrestrial communication through the Himalaya mountain range, so it is reasonable to keep South Asia and East Asia as separate regions. Telegeography provides an exact mapping of countries to regions.

Whereas Telegeopraphy reports potential and lit bandwidth at the level of the cable, it reports used bandwidth annually at the level of the region pair.[5] For instance, we observe the level of used bandwidth between North America and Europe. Used bandwidth is not directional. We observe a single used bandwidth number for the region pair, the total of used bandwidth going in both directions.[6] We observe used bandwidth between each of the seven regions. However, if there is no major cable between two regions, used bandwidth for that region pair is naturally zero. Importantly, used bandwidth reflects bandwidth between the two regions, although used

[4]Even countries with land routes between them may use subsea cables. Building subsea cables is often cheaper than building terrestrial cables and so is an attractive option even when terrestrial cables are feasible. For instance, there are a number of subsea cables running along the coast of Europe (such as between Italy and Greece).

[5]Telegeography often reports data at higher levels of aggregation and in fact refers to what we call "regions" as "sub-regions." Fortunately, Telegeophgraphy was able to provide used bandwidth data at Telegeography's subregion level.

[6]Our understanding is that used bandwidth is almost always sold symmetrically. That is, wavelengths sold in one direction come with the same number of wavelengths in the other direction.

bandwidth may carry traffic that is from further endpoints. For example, there was no major cable between Europe and South America as of 2020.[7] This means that carriers would route traffic between these regions through intermediate regions, such as Europe to North America to South America. In this case, traffic between Europe and South America would show up as used bandwidth on these intermediate routes.

Telegeography also reports international bandwidth usage for each country, that is, the total used bandwidth to other countries. This variable is reported at the level of the country, not the country pair. Similar to the interregional used bandwidth numbers, used bandwidth for a given country may reflect traffic that traverses through the country, not starting or ending there. Our model makes predictions about total used bandwidth of subsea cables so this variable is potentially useful for us. But this is one of the few Telegeography variables that includes terrestrial cables, which makes it difficult to integrate into our model. However, for island countries, all international usage is on subsea cables. We focus on island countries with substantial international bandwidth usage, of which we identify seven countries: Singapore, Japan, Taiwan, Australia, South Korea, the Phillippines, and New Zealand.[8] Section 4.1 describes how including islands as separate regions provides additional moments in our estimation.[9]

Our analysis then focuses on interregional communication (including communications between a region and an island and communications between two islands), which must traverse subsea cables. For the same reason, we restrict our sample to cables that connect separate subregions. For instance, we drop cables that connect

[7]The only viable cable in 2020 was Atlantis-2, which had a small potential capacity of 0.16 Tbps and was decommissioned in 2022.

[8]In our framework, the United Kingdom is not an island country. Significant bandwidth runs through the Channel Tunnel and is not classified as subsea, and does not appear in our data set. South Korea is not physically an island but because no South Korean communication runs through North Korea, its only land connection, South Korea is like an island for our purposes.

[9]Even if we did observe country-to-country used bandwidth, it would not be useful for us in our approach. That is because we would not know the beginning and endpoints of the usage and our model does not make predictions about that. For instance, suppose we observed used bandwidth between France and the United States in addition to used bandwidth between Europe and North America. We would know how much of that usage originated elsewhere in Europe rather than France, and because we do not model terrestrial networks (for which we lack data), our model does not make predictions about France-to-US usage. As shown below, our model makes predictions about how much data France and the US communicate (that is, with France and the US as endpoints), how much goes between the US and Europe (potentially with other endpoints) but not how goes between the US and France.

Italy to Greece or connect among Indonesian islands.

We define *markets* as pairs of regions (e.g. North America-Europe, North America-South America). Lastly, we define *path* as a set of subregion pairs for which there is a cable that connects two subregions in the market. For example, possible paths between North America and Europe include {North America-Europe; North America-South America-Europe; North America-South America-Sub-Saharan Africa-Europe; etc}. We refer to the path with just the two endpoints on it as the *direct path*. So, in the previous list, the first element is the direct path and the rest are indirect paths.

An important input into our model is the amount of capacity in each of the paths between two regions in a given year. We construct that as follows. If a cable in our data set has landing points in two markets, we denote that cable as providing capacity in the direct path between those markets. For instance, if a cable has Northern France and New Jersey as its only landing points (the Apollo cable), we denote it as connecting the regions of North America and Europe. To find the capacity on the direct path between two regions, we add up all of the capacity for the cables in that path (that is, the capacity for the cables in that market, since direct paths correspond to markets). Similarly, the capacity of the Apollo cable would contribute to any indirect path involving Europe and North America. For instance, one of the paths from Southern Africa to North America is South Africa to Europe to North America. The Apollo capacity is included as part of the capacity on the second leg of this path.

Some cables have multiple landing points. If all of the landing points fall in only two regions, we treat the cable as providing capacity on the direct path (as if it had only two landing points). If a cable has landing points in multiple regions, we denote it as providing direct connections between each of the regions. For instance, a cable that goes from East Asia to South Asia to Europe (such as the SeaWeMe cables) provides capacity on the direct path between each pair of regions. Furthermore, when a cable contributes to the direct path between two regions, we do not also record it as contributing to an indirect path. For example, SeaWeMe contributes capacity to the direct path between Europe and South Asia, the direct path between South Asia and East Asia, but not the indirect path Europe-to-South-Asia-to-East Asia. That would be double counting its capacity because SeaWeMe already contributes to the direct path between Europe to East Asia. In this sense, whether capacity appears in a market depends on which pair of regions we are

considering data traveling between. SeaWeMe contributes capacity in the Europe-South Asia market if we are considering data sent from France to India but not from France to China.

We construct other measures of cable features similarly. For instance, we construct variables such as the distance of a path and the number of cables on a path. We construct path distance by summing the market-level cable lengths for each market in a path. We construct path capacity by taking the minimum capacity of market-level capacities (alternatively, maximum or weighted average) of each market on a path. The number of cables in the path is computed as the capacity-weighted average across markets connected by the path. We take this as a measure of the level of competition on the path.

We supplement Telegeography's data with gravity data from Centre d'Etudes Prospectives et d'Informations Internationales (hereafter CEPII). It provides information used to estimate gravity equations for traditional trade such as GDP, population, geographical distance, proxies for cultural proximity, trade flows, and information on trade agreements and international relationships. We drop countries that ever had a GDP of lower than $1 billion and countries with missing GDP information, arriving at a sample of 161 countries.

Table 1 provides information on the final sample of the cables, owners, and markets included in our analysis. There are a total of 102 unique cables, with an average potential capacity of 33.8 Tbps and an average length of 10,713 km. The average construction cost for a cable is 467 million dollars. An average owner owns 2.9 cables with 3 Tbps, and is present in 4.2 markets. Our sampe has 21 unique markets. On average, a market is connected by 3.6 cables with 73.1 Tbps and 29.1 owners. Markets vary widely in their characteristics, with the largest capacity market, connecting North America and Europe, having 21 cables with a total capacity of 1,118 Tbps in 2020.

We show the overall trend in the construction and ownership of cables in Figure 2. Panel (a) shows that there is active investment throughout the sample period, with a slight rise in the number of investment episodes in the 2017-2020 period. The total capacity added in this period much larger as well, partially due to the increase in the average capacity of new cables. Panels (b)-(d) show the distributions of the number of investments, the number of cables, and the number of markets served by owner, respectively. These distributions are generally very skewed. For

example, although the average number of cables a firm owns is 2.9, but some firms own over 30 cables. Similarly, 48% of the owners serve a single market, but 10% of owners serving 10 or more markets with some firms serving 15 markets.

2.3 Motivating Facts

Table 2 shows the overall trend in international bandwidth usage. Total international bandwidth usage grew by a factor of 1000 from 2002 to 2020 as shown in the first column. Columns (2)-(4) show measures of concentration in usage at the country level. In general, the network of data trade is becoming less concentrated. In 2002, only 14 countries accounted for 90% of bandwidth usage, while by 2020 the number grew to 24 countries. We also compute the HHI and find that it falls substantially from 1076 to 719.

We now turn to patterns in the investments in the cable network, in particular, the geographical imbalance. As previously mentioned, there is a potential misallocation of cable capacity if there is a high social value of connecting under-served regions. In Figure 3, we plot total capacity and new investment episodes for the top 3 markets by total caapcity in 2020 in Panel (a) and for the bottom 3 markets (among those that have positive investment) in Panel (b).[10] We find that new investments tend to be concentrated on routes that already have high existing capacity both in terms of the number and size of investments.

Specifically, we observe large investments in the Europe-North America and East Asia-North America markets throughout the 2002-2020 period, while the North America-South America market is attracting more investments in the later period of 2014-2020. By contrast, the first South America-Africa cable with substantial capacity does not appear until 2018. These differences are potentially driven mostly by heterogeneity in demand for data exchange across markets. There may be other important factors, however, such as business stealing incentives and investment externalities across markets contributing to these patterns. Our model proposed in Section 3 allows us to distinguish these factors.

[10]There are many markets that receive no investment in our sample period. They include North America-Sub-Saharan Africa, North America - South Asia, Oceania - South America, South America - South Asia, Oceania - Sub-Saharan Africa, East Asia - South America, Europe - South America, Europe - Oceania, and Oceania - South Asia. We drop East Asia-Sub-Saharan Africa from the figure as well because even though there is entry of one cable in 2002, its capacity is not significant (1.72 Tbps).

Given the geographical concentration that we observe, we then ask whether new cable entry impacts market outcomes (potentially adding social value) and whether that impact varies by the existing cable quality. As a preliminary analysis, we run a regression of the log difference in market-level usage on whether there was a new cable entry controlling for year and market fixed effects. We include the lagged number of existing cables (or the lagged total capacity) and its interaction with the investment indicator variable.

Table 3 shows the regression results. Given that the the ready-for-service date of a new cable can fall at any time of the year, we use the growth from the year after the investment to the following year as the dependent variable in the first two columns, while using the growth from the year of the investment to the next year in the last two columns. Across all specifications, the investment variable has a positive coefficient, suggesting that new cable entry is associated with market expansion. The first column, for example, shows that new cable entry is associated with an increase in the growth rate of usage by 9.6 percentage points. We find that the interaction term between the investment indicator and the lagged number of cables (or the lagged capacity) is negative and significant. This suggests that the market expansion effect falls with the size of the existing cable network and that business stealing may be more important relative to market expansion in more established markets.[11]

Lastly, we highlight the recent growth in subsea cable investments by major content providers. Historically, cable construction was dominated by telecommunication carriers, such as AT&T, British Telecom, and NTT (of Japan), who used cables to serve the public internet. However, the last five years have witnessed a dramatic increase in investment by major content providers. Anecdotally, they are interested in connecting their data centers, but telecommunication carriers are unwilling to build such specialized cables. Content providers may lease some bandwidth on their cables to carriers and the cables often have telecom coinvestors, but these cables are oriented, at least in part, towards the private communication needs of content providers. We observe four content providers with substantial ownership of subsea cables: Google, Facebook, Amazon, and Microsoft. Figure 4 shows the

[11]The general point that new cables improve internet performance is corroborated in many settings. For instance, Fanou et al. (2020) use internet measurement techniques to document reduced latency between South America and Africa after the 2018 arrival of the new cable connecting those continents.

share of capacity invested by these content providers for four time periods we consider. Since Google's first cable in 2010 linking North America and East Asia, even up to 2017, the content providers' overall share remains very small (below 3% in 2013-2017). Their investment grows at a remarkable rate in the 2018-2020 period, accounting for 73% of total capacity in the Europe-North America market, for example.

3 A Model of Data Flows and Cable Investment

3.1 Data Flows and Cable Usage

In this section, we present an estimable model of how country-to-country demand for data maps into flows between regions that we observe in our data. In our model, countries are indexed by $c = 1, \ldots, C$ and each country pair ck has demand for data from each other. Time is indexed by t. Data are measured in Gigabits per second. Every pair of countries c and k has a potential demand for data from each other of $\overline{M}$. The share of data transmitted between countries is a function of their demand for data, which we capture with demographic and country characteristics, and the quality of the connection between the two countries, which is a function of cable characteristics. The total data transmitted between the two countries is:

$$d_{ckt} = \frac{\exp(x_{ckt}\theta^d + v_{ckt})}{1 + \exp\left(x_{ckt}\theta^d + v_{ckt}\right)}\overline{M} \tag{1}$$

where x_{ckt} captures time trends and country and country-pair characteristics such as the population and GDP of the two countries The parameters θ^d are to be estimated. The variable v_{ckt} captures the quality of the connection between the two countries, further described below. The share of potential demand that is not realized could be packets that are lost in internet transit. It could also be information that consumers choose never to search for in the first place because consumers are aware of internet quality or because the cost of exchanging data is not worthwhile given the consumer's income, distance, language, and other issues.

We group countries into subregions $r = 1, \ldots, R$ and subregions into regions $g = 1, \ldots, G$. The subregions form a partition of the countries, and the regions form a partition of the subregions. As described above, we choose regions so that

communication between countries in different regions must traverse subsea cables. We use subregions to accommodate the seven island countries for which we utilize total international bandwidth data. We divide regions that contain island countries into subregions, with one subregion for each island country (these subregions contain a single country) and one subregion for the remaining countries. Regions with no island countries contain a single subregion that contains all of the countries in the region. Our model makes predictions at the subregion level that we add up to form region-to-region predictions of usage, which is the level we observe usage. We denote pairs of subregions as *markets* indexed by $m = 1, \ldots, M$. For example, North America-Europe is one market and North America-East Asia is another market. Markets are non-directional so North America-Europe is the same market as Europe-North America.

The network of cables dictates a set of paths that an internet packet might travel over. Between any two countries c and k, there exists P_{ckt} *paths* indexed by $p = 1, \ldots, P_{ckt}$. Each path is a set of subregion pairs for which there is an active cable. For instance, three possible paths from North America to Europe include {North America-Europe; North America-South America $\rightarrow$ South America-Europe; North America-Sub-Saharan Africa$\rightarrow$ Sub-Saharan Africa-Europe}. Internet cables are bidirectional so $P_{ckt} = P_{kct}$ and any sequence of markets that connect countries c to k also connects k to c. The set of paths can evolve as new cables are constructed in markets for the first time. Pairs of countries in the same respective sub-regions have the same set of paths available to them. For example, the US and Spain have the same paths between them as Canada and France.

The way in which data travels over the internet is a complex function of a number of variables. In general, routing behavior depends on traffic, so internet routers will automatically direct packets of information to the least crowded routes. Even packets from the same overall stream, for instance, a video or email message, may take very different routes to their destination. We use a simple reduced-form approximation for this process.

Each path is characterized by an *attractiveness* δ_{pckt}. Attractiveness δ_{pckt} depends on cable features such as path distance and bandwidth associated with each path as well as the level of competition between cable owners in the markets that make

up the path.[12] The share of data going from c to k on path p at time t is:

$$s_{pckt} = \frac{\exp\left(\delta_{pckt}\right)}{\sum_{z=1}^{P_{ckt}} \exp\left(\delta_{zckt}\right)} \tag{2}$$

We parameterize the attractiveness of the path δ_{pckt} to be a linear function of path characteristics as follows:

$$\delta_{pckt} = Z_{pkt}\theta^{\delta}$$

where Z_{pkt} includes a constant and path characteristics such as the length, the potential capacity, and the number of cables of path p. We also use HHI for ownership share on that path. We construct potential capacity by summing over potential capacity for cables that have landing points in both subregions containing countries c and k.

We let the quality of the connection between two countries equal the inclusive value of logit choice among cables:

$$v_{ckt} = \ln\left(\sum_{p=1}^{P_{ckt}} \exp\left(\delta_{pckt}\right)\right).$$

Thus, $s_{pckt} = \exp\left(\delta_{pckt}\right) / \exp(v_{ckt})$. The amount of data traveling on a given path between two countries is:

$$\hat{d}_{pckt} = s_{pckt} d_{ckt} = \frac{\exp\left(\delta_{pckt} + x_{ckt}\theta^{d}\right)}{1 + \sum_{z=1}^{P_{ckt}} \exp\left(\delta_{zckt} + x_{ckt}\theta^{d}\right)} \overline{M}.$$

In our model, more paths between two countries (higher P_{ckt}) and the increased quality of those paths (higher δ_{pckt}, such as because of higher capacity) lead those countries to communicate more data (more share drawn from the outside option). In this way, the model captures that the quality of connectivity affects the quantity of data transmission. We do not model network congestion explicitly because congestion, in the form of reduced consumer service, is typically realized only for

[12] Local carriers and content providers negotiate with cable owners to obtain passage. Breakdowns in negotiations mean the cable will not accept that traffic, akin to an engineering constraint on how packets can be routed. In this sense, more competitive ownership structures can lead to more traffic on a path.

short periods of time and we have annual data on usage, but the model captures congestion in the sense that better connectivity leads to more network use.

Let $\mathcal{L}_{pckt}$ be the set of subregion pairs associated with path p between country pair ck. The total demand between countries c and k traveling on subregion pair l in period t can then be obtained by summing $\hat{d}_{pckt}$ over all paths while assigning it to all associated subregion pairs as follows.

$$\tilde{d}_{lckt} = \sum_{p=1}^{P_{ckt}} \hat{d}_{pckt} \mathbb{1}\{l \in \mathcal{L}_{pckt}\}.$$

In estimation, we would like to match the predicted data flows to the observed bandwidth usage in the data. A challenge is that paths are defined to be pairs of subregions, but we observe used bandwidth between markets, i.e., pairs of regions. Hence, we need to sum over subregion pairs to obtain total data flows in a market. Let l index pairs of subregions, and let $\mathcal{S}_m$ be the set of subregion pairs contained in market m. For instance, the region *East Asia* has three subregions, *Japan*, *Korea* and *Rest of East Asia*. The region *North America* has only one subregion. For m=*East Asia-North America*, $\mathcal{S}_m$ contain three subregion pairs, *Japan-North America*, *Korea-North America* and *Rest of East Asia-North America*.

Thus, total data in market m in period t is:

$$D_{mt} = \sum_{l \in \mathcal{S}_m} \left(\sum_{c=1}^{C} \sum_{k=1}^{c-1} \tilde{d}_{lckt} \right) \tag{3}$$

We denote total usage for country c in period t to be Q_{ct}. In our model, this equals:

$$Q_{ct} = \sum_{k=1}^{C} d_{ckt} \left(\sum_{p=1}^{P_{ck}} s_{pckt} \right).$$

This quantity is also useful, as bandwidth usage is available at the country level in addition to the market level in our data.

3.2 Cable Investment

In our model, firms $j = 1, \ldots, J$ simultaneously choose whether to invest in a new cable in each market $m = 1, \ldots, M$. Firms profit from their ownership share of

cables and the amount of data that traverses those cables. Each market is characterized by an amount of capacity available on that market K_{mt}. The amount of capacity owned by firm j is denoted by k_{mjt}.

In each period, a new line may be constructed in a market. Firms observe the capacity of the potential cable, $\tilde{k}_{mt}$. Then, each firm makes a binary choice $a_{jmt} \in \{0, 1\}$ in each market whether to build the available line. Let $\boldsymbol{a}_{jt} = (a_{jmt})_m$ denote a vector of firm j's investment choices in all markets. If multiple firms choose to build in a given market, they become joint owners of the line, evenly splitting ownership of the line. Owners split the cost of the line $\omega_{mt}(\tilde{k}_{mt})$ evenly. Each owner also receives a firm-specific investment cost shock (ν_{jmt}). The total cost of investment for market m, is then given as follows:

$$c_{mt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \tau\right) + a_{jmt}\nu_{jmt} = \frac{a_{jmt}}{\sum_{k=1}^{J} a_{kmt}}\omega_{mt}(\tilde{k}_{mt}) + a_{jmt}\nu_{jmt}.$$

In our framework based on Pakes et al. (2015), the value of ν_{jmt} can be common knowledge across firms or privately observed.

The profit for firm j is given by

$$\bar{\pi}_{jt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma, \tau\right) = \sum_{m \in \mathcal{M}} \left\{r_m(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma) - \left(c_{mt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \tau\right) + a_{jmt}\nu_{jmt}\right)\right\}.$$

where $r_m(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma)$ is the revenue in market m from choosing $\boldsymbol{a}_{jt}$ (i.e. the total profit before paying the investment cost) when the competitors choose $\boldsymbol{a}_{-jt}$. X_{jt} is the set of all exogenous variables that affect the profit, including firm j's capacities and total capacities in each market as well as all variables that affect D_{mt}. Firms have measurement error over elements of $r(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma)$, which is represented by the additive term ε_{ajmt}. Thus, letting $\mathcal{J}_{jt}$ be the set of information firm j uses in forming expectations about the revenue to choosing $\boldsymbol{a}_{jt}$, firm j perceives the profit from investment in mt as:

$$
\begin{aligned}
&\pi_{jt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma, \tau\right) \\
&\quad = E\left[\sum_{m \in \mathcal{M}} \left\{r_m\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma\right) - \left(c_{mt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \tau\right) + a_{jmt}\nu_{jmt}\right)\right\} \middle| \mathcal{J}_{jt}\right] \\
&\quad = \sum_{m \in \mathcal{M}} \left\{r_m\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma\right) + \varepsilon_{ajmt} - \left(c_{mt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}\,; \tau\right) + a_{jmt}\nu_{jmt}\right)\right\}.
\end{aligned}
$$

Firms play a Nash equilibrium, simultaneously solving:

$$
\max_{\boldsymbol{a}_{jt}} \pi_{jt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma, \tau\right).
$$

Let $\tilde{\boldsymbol{a}}_{jt}^{m,1}$ be the vector $\boldsymbol{a}_{jt}$ where the mth element is set to 1. We similarly define $\tilde{\boldsymbol{a}}_{jt}^{m,0}$. That is, $\tilde{\boldsymbol{a}}_{jt}^{m,1} = (a_{j1t}, ..., a_{jm-1t}, 1, a_{jm+1t}, ..., a_{jMt})$ and $\tilde{\boldsymbol{a}}_{jt}^{m,0} = (a_{j1t}, ..., a_{jm-1t}, 0, a_{jm+1t}, ..., a_{jMt})$. The incremental change in the revenue in market z to j from investing in market m is:

$$
\Delta r_{mz}(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}) = r_z(\tilde{\boldsymbol{a}}_{jt}^{m,1}, \boldsymbol{a}_{-jt}, X_{jt}) - r_z(\tilde{\boldsymbol{a}}_{jt}^{m,0}, \boldsymbol{a}_{-jt}, X_{jt}).
$$

Also, we define $\Delta\varepsilon_{jmt} = \varepsilon_{1jmt} - \varepsilon_{0jmt}$. Revealed preference tells us that if firm j invests in market m, t, it must be that:

$$
\sum_{z \in \mathcal{M}} \Delta r_{mz}(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma) + \Delta\varepsilon_{jmt} - \left(c_{mt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}\,; \tau\right) + a_{jmt}\nu_{jmt}\right) \geq 0 \quad (4)
$$

We take the expectation of equation (4) conditional on the firm's information set $\mathcal{J}_{jt}$ to obtain

$$
E\left[\sum_{z \in \mathcal{M}} \Delta r_{mz}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma\right) - \left(c_{mt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}\,; \tau\right) + a_{jmt}\nu_{jmt}\right) \middle| a_{jmt} = 1, \mathcal{J}_{jt}\right] \geq 0
$$

where $E[\Delta\varepsilon_{jmt}|\mathcal{J}_{jt}, a_{jmt}] = 0$ follows from the rational expectations assumption. Suppose that $Z_{jmt} \subsetneq \mathcal{J}_{jt}$ is the set of variables in the firm's information set that are exogenous in the sense that they do not respond to changes in a_{jt} and a_{-jt}. Then,

we can apply the law of iterated expectations to derive:

$$E\left[\sum_{z\in\mathcal{M}}\Delta r_{mz}(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt};\gamma)-\left(c_{mt}\left(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt}\,;\tau\right)+a_{jmt}\nu_{jmt}\right)\Big|a_{jmt}=1,Z_{jmt}\right]\geq 0. \tag{5}$$

Similarly, we can define a separate moment for the case of no investment:

$$E\left[\sum_{z\in\mathcal{M}}\Delta r_{mz}(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt};\gamma)-\left(c_{mt}\left(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt}\,;\tau\right)+a_{jmt}\nu_{jmt}\right)\Big|a_{jmt}=0,Z_{jmt}\right]\leq 0. \tag{6}$$

It remains to specify the revenue function $r_m()$ and the investment cost function $\omega_{mt}()$. We specify the revenue function as:

$$r_m(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt};\gamma)=\gamma D_{mt}\frac{k_{jmt}\left(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt}\right)}{K_{mt}\left(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt}\right)}=\gamma q_{jmt} \tag{7}$$

where γ is the parameter of our interest and q_{jmt} denotes the quantity served by firm j in mt. The total data transmitted in market m, denoted $D_{mt}(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt})$, depends not only on the total capacity in market m, but also on the capacities in all other markets, as well as interregional demand for data flows across all markets. The variables K_{mt} and k_{jmt} are determined based on investment choices as follows:

$$K_{mt}\left(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt}\right)=K_{mt-1}+\tilde{k}_{mt}\mathbf{1}\left\{\sum_{j=1}^{J}a_{jmt}\geq 1\right\}$$

$$k_{jmt}\left(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt}\right)=k_{jmt-1}+\frac{a_{jmt}}{a_{jmt}+\sum_{-j}a_{-jmt}}\tilde{k}_{mt}$$

We specify the investment cost as a linear function of the cable's capacity and length in logs.

Equations (5) and (6) still cannot be applied to the data as they depend on unobserved structural errors (ν_{jmt}), which causes the well-known selection issue. That is, conditional on choices, these errors are not mean zero. To move forward, we assume $\nu_{jmt}=0$ for all j,m,t, obtaining the following inequality from equation

(5) :

$$E\left[\sum_{z\in\mathcal{M}}\Delta r_{mz}(\boldsymbol{a}_{jt,-jt}\,,X_{jt};\gamma)-c_{mt}\left(\boldsymbol{a}_{jt},\boldsymbol{a}_{-jt},X_{jt}\,;\tau\right)\Big|a_{jmt}=1,Z_{jmt}\right]\geq 0. \qquad (8)$$

We can define the analogous inequality for the case of no investment.

4 Estimation

4.1 Data Flows and Cable Usage

We briefly discuss the intuition behind our identification of the demand-side model. Identification is challenging because we wish to infer both country-to-country demand and cable quality from observed data flows. Taking the simplest example possible of two regions connected by a single path for one period, we would observe only one data flow (that is, only one observation of the dependent variable). Naturally, we cannot infer both region-to-region demand and path quality in this case. Even as the number of observations increases, for instance, because of an increase in the number of regions or time periods, a fully non-parametric treatment (that is, non-parametric treatment of $x_{ckt}\theta^d$ that varied across all region pairs- time combinations as well as a non-parametric treatment of paths (that is, treating δ_{pckt} non-parametrically, even assuming it is constant across country pairs in the same region pairs) would be infeasible. A further challenge is that there are multiple countries per region and we wish to identify the effect of some country-to-country variables.

We make reasonable functional form assumptions. We allow for region-to-region fixed effects that are constant over time and allow region-to-region demand to be further influenced by demographic variables and some match variables (particularly, distance). This implies that as the number of regions grows, the number of observations grows relative to the number of parameters because the demand that a region has with one region cannot grow in a way that is totally disconnected to the demand with another region. Moments drawn from island countries contribute to precision over country-level variables. Cable quality parameters in δ_{pckt} are determined by the level of usage observed in markets conditional on country-to-country demand.

In estimation, we match D_{mt} and Q_{ct} to the observed used bandwidth in the data. We denote the predicted data flows for market m and time t given parameter θ as $\tilde{D}_{mt}(\theta)$ where $\theta = [\theta^d, \theta^\delta]$ includes parameters that govern demand between countries (θ^d) and parameters that govern how demand gets allocated across the cable network (θ^δ). Let $\overline{D}_{m.}(\theta)$ is the average over time of $\tilde{D}_{mt}(\theta)$, and $\overline{D}_{.t}(\theta)$ is the average over m. Similarly, we denote the predicted international data flow for country c and time t given parameter θ as $\tilde{Q}_{ct}(\theta)$. We let $\overline{Q}_{.t}(\theta)$ denote the average of $\tilde{Q}_{ct}(\theta)$ over island countries and $\overline{Q}_{.t}(\theta)$ the average over time.

Given the parameter values, we compute the data flows for each market and time period based on equations (3). We construct the set of moments by stacking the average predicted data flows at the market level and year level:

$$\Gamma(\theta) = \left[\overline{Q}_{c.}(\theta), \overline{Q}_{.t}(\theta), \overline{D}_{m.}(\theta), \overline{D}_{.t}(\theta)\right].$$

Let Z be the set of instruments such that $E\left[Z'\left(\Gamma^d - \Gamma(\theta)\right)\right] = 0$ where Γ^d is the set of moments based on the observed bandwidth usage in the data. We search for the parameter vector that minimizes the weighted distance between the predicted and observed moments given as

$$f(\theta) = \left(\Gamma^d - \Gamma(\theta)\right)' Z' W^{-1} Z \left(\Gamma^d - \Gamma(\theta)\right)$$

where W is the consistent estimate of $E[Z'\left(\Gamma^d - \Gamma(\theta)\right)\left(\Gamma^d - \Gamma(\theta)\right)' Z]$. We use constants and time trends as instruments and set $\overline{M} = 4000$.

Panels A and B of Table 5 presents the demand-side estimates. The estimates suggest that a one-percent increase in the sum of the GDP of two countries involved in trade would result in an approximately 4 percent increase in trade in data, and a one-percent increase in the sum of the population would result in 1.6 percent decrease in trade in data. Thus, a one-percent increase in both GDP and population (an increase in GDP per capita) leads to an increase in trade in data. Further, consistent with our expectations, the path length has a negative effect, while the path capacity and number of cables have positive effects on the amount of data flows allocated to the path. One way to interpret the number of cables is as a measure of market structure among cable owners on the path, in which case the positive effect would reflect the enhanced competition when there are more cables. Another interpretation is about redundancy, as carriers sometimes value having

more options on a path in case one cable malfunctions.

4.2 Cable Investment

We derive the following inequalities by interacting inequalities in (8) with positive valued functions of Z_{jmt}, denoted as $h(Z_{jmt})$.

$$\frac{1}{N^{a_{jmt}=1}} \sum_{\{j,m,t|a_{jmt}=1\}} \left(\sum_{z \in \mathcal{M}} \Delta r_{mz}(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}; \gamma) - c_{mt}\left(\boldsymbol{a}_{jt}, \boldsymbol{a}_{-jt}, X_{jt}\,; \tau\right) \right) h(Z_{jmt}) \geq 0 \tag{9}$$

where $N^{a_{jmt}=1}$ is the number of observations (firm-market-year combinations) for which the firm chooses to invest. We can define an analogous inequality for the case of no investment.

We focus on recovering the revenue function while calibrating the cost of investment using the construction data for each cable.[13] We estimate the construction cost by regressing the logged cost on a constant, the length, and the capacity of the cable in logs using only cables that connect one market (see Appendix Table A1 for detailed results). We use predicted values from this regression as the investment costs.

A final step that is necessary to apply the inequalities is specifying counterfactual outcomes. That is, given that we observe investment by a firm in a certain market, what would happen to the market if the firm decides not to invest, holding the strategies of its rivals constant? In the case of investment, we assume that the cable of the same capacity as observed in the data would be built if at least one other firm invested in that market and year. Otherwise, we assume that no new cable would be added. In the case of no investment, if we observe a cable built in that year and market, we assume that there will be a cable of the same capacity in the counterfactual scenario. If there is no cable built, we use the predicted capacity from a regression of capacity on time and market fixed effects as the counterfactual capacity. The results from this capacity regression are presented in Appendix Table A2.

We briefly discuss the intuition for identification of the supply-side model. Recall

[13] Another example of a paper using data on observed transaction costs to study investment in a capital-intensive industry is Jeon (2022).

that q_{jmt} denotes quantity served by firm j in mt. Denoting Δq_{jmzt} as the incremental change in quantity served in market z induced by investing in market m and simplifying the notation, equation (9) takes the following form:

$$\frac{1}{N^{a_{jmt}=1}} \sum_{\{j,m,t|a_{jmt}=1\}} \left(\gamma \sum_{z \in \mathcal{M}} \Delta q_{jmzt} - c_{jmt} \right) \geq 0 \tag{10}$$

when we use a constant as the instrument function. The sum of quantity changes across all markets ($\sum_{z \in \mathcal{M}} \Delta q_{mzt}$) is on average positive in γ, which means that the inequality will slope upward, yielding a lower bound for γ. An increase in the quantity term or a decrease in the cost term will shift this line to the left, thereby resulting in a smaller lower bound. Intuitively, observing firms investing only when the cost of investment is low, holding everything else constant, for example, implies that the gain from investing must be relatively small. Similarly, observing firms only when the quantity gain is large implies that the profit from the incremental quantity gain from the investment must be relatively low. Observations in which firms choose not to invest yield an upper bound following the same logic. If we observe firms choosing not to invest when the cost of investment is high (or the quantity gain is low), we would estimate a larger upper bound for γ.

We allow the markup parameter γ to vary by firm type (content providers vs. non-content providers) and across four time periods. Our inference method is based on Cox and Shi (2023) with detailed steps provided in Appendix B. Panel C of Table 5 presents the estimates of γ. In the 2002-2007 period, we recover a 95% confidence interval of (5230, 7760) for non-content-providers, which implies that an increase in a firm's quantity ($\sum_{m \in \mathcal{M}} D_{mt} \frac{k_{jmt}}{K_{mt}}$) by 1 Tbps would lead to a profit increase ranging from 5.2 to 7.8 billion dollars. By 2018-2020, this number decreases to 11 to 12 million dollars. This steep decline in the markup is not surprising, considering the similar steep decline in bandwidth prices.[14]

Our estimates of the markup parameter for content providers are substantially lower than the estimates for non-content providers. They range from 5.1 to 29.4

[14]For example, the compound annual growth rate from 2016 to 2019 for the median monthly lease price for 10 Gbps wavelength ranges from -12 to -45% on Europe-Asia routes, -4 to -15% for Transatlantic routes, and -24 to -36% U.S. to Latian America routes according to the sample of prices available in our Telegeography data. As a comparison, we compute the annual growth rate of the markup term using the middle point in each of the time periods and the lower bounds. The rate is approximately -10% from 2002-2007 to 2018-2020 and -17% from 2013-2017 to 2018-2020.

million dollars per Tbps in 2002-2007 and 4.6 to 12.3 million dollars per Tbps in 2018-2020. This is driven by the fact that content providers tend to invest in periods in which the quantity gain from investing is higher than non-content providers, even though they tend to invest when the investment cost is also higher as shown in Figure 5. Similarly, the quantity increase content providers would have enjoyed from investing when they chose to withhold is higher. These patterns in the data lead to smaller lower and upper bounds of γ. This finding is consistent with our understanding of the industry. Many industry experts report that content providers have distinct incentives from traditional telecommunication carriers in that they are particularly concerned about connecting their data centers and serving their own growing demand (for example, Satarinano, 2019). Furthermore, unlike telecommunication carriers, they are not allowed to sell capacity directly but instead use it for their own consumption or swap capacity with other firms.

We highlight the network nature of the market by examining how much of the data transmissions on cables in a given market can be attributed to actual endpoint-to-endpoint usage. For example, Table 4 shows that 84% of the usage in cables in the Europe-North America market in 2020 can be attributed to data exchanges between Europe and North America, while the rest is due to usage in other markets such as East Asia-North America and North America-South Asia.

5 Counterfactuals

5.1 Understanding the Growth in Trade in Data

A limitation in the observed bandwidth usage data is that it does not capture endpoint-to-endpoint usage, because data traveling between two countries can take multiple routes determined by various factors. Our structural model allows us to map observed data flows to country-to-country data exchanges. That is, given our estimates of the demand parameters, we can infer country-to-country demand for data. In addition, our model allows us to study which factors have contributed to the growth in trade in data. We decompose the observed growth into the part driven by (a) changes in demand for data exchange, and that driven by (b) improvements in the quality of cable connections. The former is captured by changes in x_{ckt} including demographic characteristics and time trends and the latter by changes in

v_{ckt} in the model.

We first compute the total global data flows based on our model from 2002 to 2020, which is shown as the blue line in Figure 6. To isolate the effect of the changes in the quality of the cable connection, We compute the total data flows with x_{ckt} fixed at the 2020 level and setting v_{ckt} to its observed level for each year. Similarly, to isolate the effect of changes in x_{ckt}, we compute the total data flows with v_{ckt} fixed at the 2020 level while setting x_{ckt} to its observed level each year.

Results appear in Figure 6. Because the two effects interact, the sum of each isolated effect is often less than the total effect. We find that in 2002, for example, enhancing the cable network to the 2020 level while fixing everything else at the 2002 level would not lead to a substantial increase in data transmissions. Endowing the world with the 2020 level of x_{ckt} while fixing the cable network at the 2002 level, however, would bring the data flows up to approximately one-half of the 2020 level. In the later part of the sample period, however, the demographic characteristics (and the time trend) and the quality of cable connections have comparable contributions to the growth in data flows. For example, in 2014, improving the cable connections to the 2020 level while fixing x_{ckt} at the observed level would double the data flows from the observed level or bring it up to 36% of the 2020 level.

5.2 Inefficiencies in the Subsea Cable Network

The high level of geographical concentration in the cable network suggests the possibility of misallocation of cables to markets that receive relatively low social values from these cables. We characterize these inefficiencies in the cable network and quantify the business stealing effects based on our estimates. For each market, we consider constructing a new cable by a de novo entrant with no cable. We assume that the capacity of the new cable is given in the same way we construct counterfactual outcomes in the estimation based on the capacity regression described in Section 4.2. We also assume that all other firms, or incumbents, take actions observed in the data. Note that we are considering the effect of a new cable rather than a new owner joining an existing cable.

For each counterfactual scenario, we compute the effect of the new cable on incumbent profits (the sum of profits over all firms in our data). Unlike in other

settings with independent markets, an investment in this setting affects incumbents' profits not only in the market where the investment happens but also in all other markets due to the fact that markets are interconnected. Our model allows us to decompose the effect of a new cable into these two effects.

The direction of these effects is theoretically ambiguous. A new cable would transfer market shares from incumbents to the entrants in the investment market. Furthermore, an increase in the capacity in that market would divert data flows away from the other markets, lowering profits in these markets as well. At the same time, the increased capacity would increase demand for data flows, resulting in market expansion, typically in multiple markets.

We focus on the year 2015 and separately compute the effect of adding a cable to each market. For each cable introduction, We compute changes in the incumbent profit for the market in which the investment happens and all other markets separately. We also report the effects for two markets that involve North America and two markets that involve Sub-Saharan Africa. In contrast to the North American markets that have received repeated, large investments throughout the sample period, Sub-Saharan Africa has generally been considered underserved. For example, Sub-Saharan Africa-East Asia has no major cable connection as of 2023, and Sub-Saharan Africa-South America received the first substantial investment in 2018. We also report the average of the effects across all markets.

Table 6 shows the results computed based on the lower and upper bounds of γ. We find that incumbents would collectively lose 81 to 234 million dollars in profit in the market in which a new cable arrives on average. This suggests that the business stealing effect dominates the market expansion effect for incumbents in the investment market. In contrast, in the rest of the world, incumbents would receive a profit gain of 13 to 39 million dollars. The entrant would enjoy an average variable profit of 118 to 341 million dollars from the entry.

The breakdown by market shows that the new investment would lower the total incumbent profit in the two markets involving North America. The effect is negative both for and outside the investment market, with much of the negative effect coming from the investment market. This suggests that there are substantial transfers from the incumbents to the entrant and that market expansion is limited. By contrast, in both of the Sub-Saharan markets, incumbent profits would increase substantially following the new cable entry. The revenue for the entrant also varies

widely across markets. It ranges from 6 to 52 million dollars for the Sub-Saharan markets while ranging from 76 to 780 dollars for the North American markets.

In order to compare the effect of business stealing across markets, we calculate the ratio of the negative value of the total change in the incumbent profit (in the third column) to the entrant revenue (in the fourth column). As discussed above, the North American markets have strong business stealing effects. The business stealing effect is negative in the Sub-Saharan African markets as the incumbent profit increases from increased demand following the cable entry.

The results highlight that private incentives do not align with what is socially optimal. Markets in which cable investment would yield high social gains do not necessarily guarantee high private gains. This suggests that there is a substantial misallocation of investment, and targeted subsidies could lead to a more efficient network. Because of the global nature of the externality, individual national governments are not sufficiently incentivized to address the difference between social and private gain. A national government or even a regional consortium of governments could subsidize cable construction to their area, but their benefit from doing so would presumably not account for the benefits that the cable would create in other regions.

6 Conclusion

Undersea internet cables are a critical piece of communication infrastructure, carrying more than 99% of international data traffic and underpinning modern global commerce. There is little economic research on this market despite massive attention to the digitization and e-commerce that it enables. Our paper provides two main contributions.

First, we treat global traffic on the cables as *international trade in data* and study the evolution of trade in data with trade in goods. We utilize a new data set on cable construction and use to estimate a model of country-to-country trade in data, akin to a gravity equation model well-known in the literature on international trade. We decompose the growth of trade in data into an amount due to increases in demand and an amount due to investment in the network. We find that cable construction was a substantial contributor, on par with increases in demand.

Second, we develop and estimate a model of cable construction. Our model

accounts for how constructing a cable between any two locations can affect global traffic flows according to the logic of internet transmission. In addition, our model endogenizes the formation of consortiums of investors, which are frequently observed in practice. We also allow for content providers, such as Google and Microsoft, to have different monetization than telecommunication carriers, such as AT&T and British Telecom, as the recent growth of content-provider investment has been an important development in the industry.

Our model of investment allows us to contrast private investment incentives with socially optimal investment incentives and we find that they deviate substantially. The importance of business stealing is consistent with our descriptive evidence, where we document repeated investment in busy routes and practically nonexistent investment in routes with lower demand. Importantly, the global internet backbone industry is characterized by a very low level of economic regulation. Our results highlight how appropriate subsidies could lead the industry towards a more efficient global network.

Overall, we provide a new model of undersea internet cable construction and use and estimate the model on new data on this important market. Our results emphasize the importance of cable construction in global data flows and the likelihood that the market delivers inefficient outcomes. Future research on adjacent markets, such as the growth of data centers, appears valuable.

References

Aguirregabiria, V. and Ho, C.-Y. (2010). A dynamic game of airline network competition: Hub-and-spoke networks and entry deterrence. *International Journal of Industrial Organization*, 28:377–382.

Allen, T. and Arkolakis, C. (2022). The welfare effects of transportation infrastructureimprovements. *Review of Economic Studies*, 89:2911–2957.

Andrews, D. W. K. and Barwick, P. J. (2012). Inference for parameters defined by moment inequalities: A recommended moment selection procedure. *Econometrica*, 80(6):2805–2826.

Blum, B. S. and Goldfarb, A. (2006). Does the internet defy the law of gravity? *Journal of International Economics*, 70(2):384–405.

Brancaccio, G., Kalouptsidi, M., and Papgeorgiou, T. (2020). Geography, transportation, and endogenous trade costs. *Econometrica*, 88:657–691.

Canay, I. A., Illanes, G., and Velez, A. (2023). A User's Guide for Inference in Models Defined by Moment Inequalities.

Caoui, E. H. and Steck, A. (2023). Network diversity, market entry, and the global internet backbone. *SSRN Electronic Journal*.

Clark, D. D. (2018). *Designing an Internet*. MIT Press, Cambridge, MA.

Cox, G. and Shi, X. (2023). Simple adaptive size-exact testing for full-vector and subvector inference in moment inequality models. *Review of Economic Studies*, 90:201–228.

CRS (2022). Undersea Telecommunication Cables: Technology Overview and Issues for Congress. Technical report.

Dickstein, M. J. and Morales, E. (2018). What do Exporters Know?*. *The Quarterly Journal of Economics*, 133(4):1753–1801.

Eizenberg, A. (2014). Upstream innovation and product variety in the United States home PC market. *The Review of Economic Studies*, 81:1003–1045.

Elliott, J. T., Houngbonon, G. V., Ivaldi, M., and Scott, P. T. (2023). Market structure, investment, and technical efficiencies in mobile telecommunications. Unpublished manuscript, New York University.

Fan, Y. and Yang, C. (2022). Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety.

Fanou, R., Huffaker, B., Moke, R., and Claffy, K. (2020). Unintended consequences: Effects of submarine cable deployment on internet routing. Unpublished manuscript, Center for Applied Internet Data Analysis.

Fréchette, G. R., Lizzeri, A., and Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review*, 109:2954–2992.

Greenstein, S. (2015). *How the Internet Became Commercial: Innovation, Privatization, and the Birth of a New Network*. Princeton University Press.

Hjort, J. and Poulsen, J. (2019). The arrival of fast internet and employment in africa. *American Economic Review*, 109:1032–1079.

Holmes, T. J. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica*, 79:253–302.

Houde, J.-F., Newberry, P., and Seim, K. (2022). Nexus tax laws and economies of density in E-commerce: A study of Amazon's fulfillment center network. *Econometrica*, In Press.

Jeon, J. (2022). Learning and investment under demand uncertainty in container shipping. *RAND Journal of Economics*, 53:226–259.

Jia, P. (2008). What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica*, 76:1263–1316.

Kline, B., Pakes, A., and Tamer, E. (2021). Moment inequalities and partial identification in industrial organization. In Ho, K., Hortaçsu, A., and Lizzeri, A., editors, *Handbook of Industrial Organization*, volume 4, pages 345–431. Elsevier.

Li, S., Mazur, J., Park, Y., Roberts, J., Sweeting, A., and Zhang, J. (2022). Repositioning and market power after airline mergers. *RAND Journal of Economics*, 53:166–199.

Pakes, A., Porter, J., Ho, K., and Ishii, J. (2015). Moment inequalities and their application. *Econometrica*, 83:315–334.

Rysman, M. (2016). Empirics of business data services. Appendix B of *Business Data Services Federal Notice of Proposed Rulemaking*, FCC 16-54.

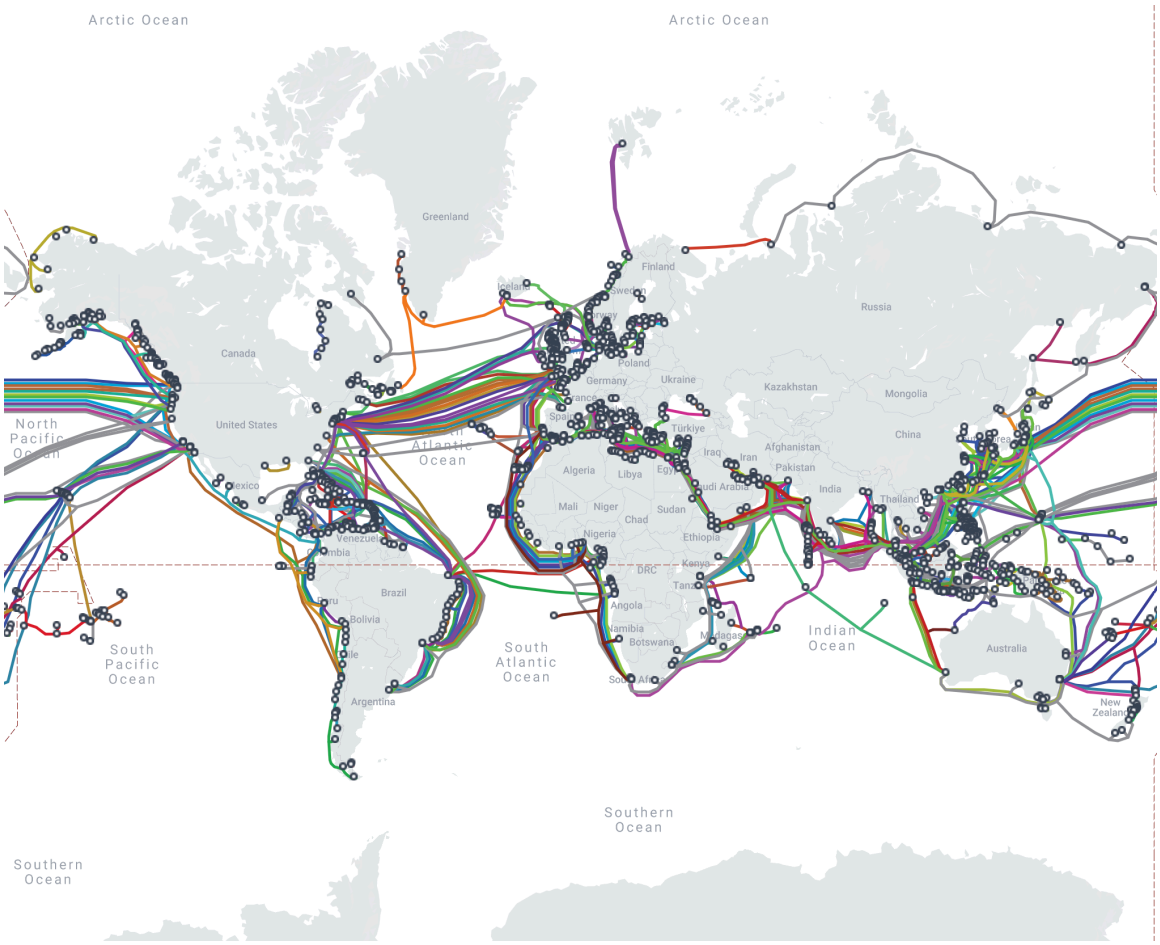
Satarinano, A. (2019). How the internet travels the ocean. *New York Times*, March 10, 2019.

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

Wollmann, T. G. (2018). Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles. *American Economic Review*, 108(6):1364–1406.

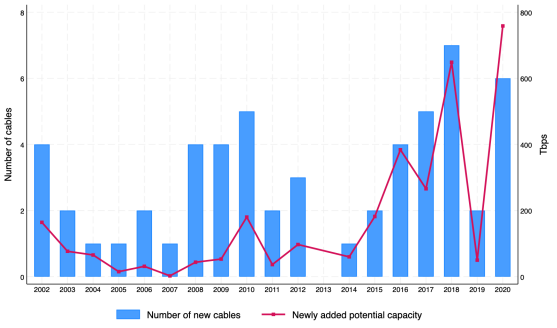
Tables and Figures

Figure 1: Map of subsea cables in 2023

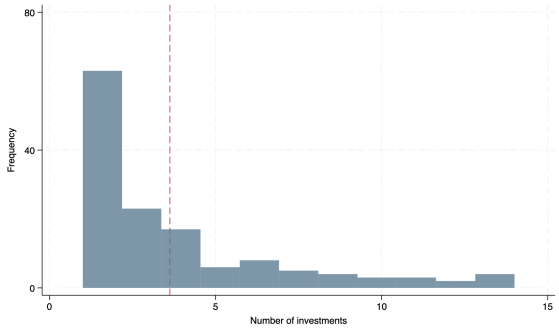


Notes: This map shows undersea telecommunications cables as of October 5, 2023. Source: TeleGeography, Submarine Cable Map, www.submarinecablemap.com

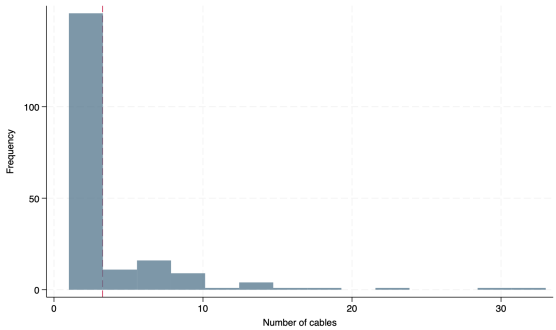
Figure 2: Subsea cable investment and ownership



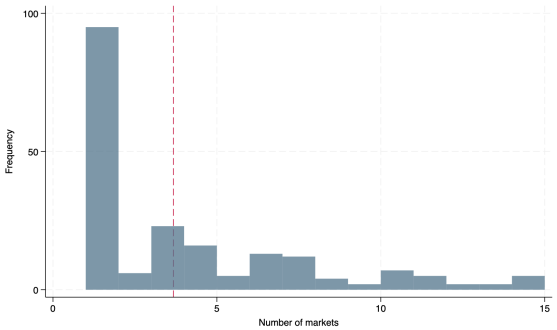
(a) Total investment



(b) Number of investments by firm



(c) Number of cables by firm



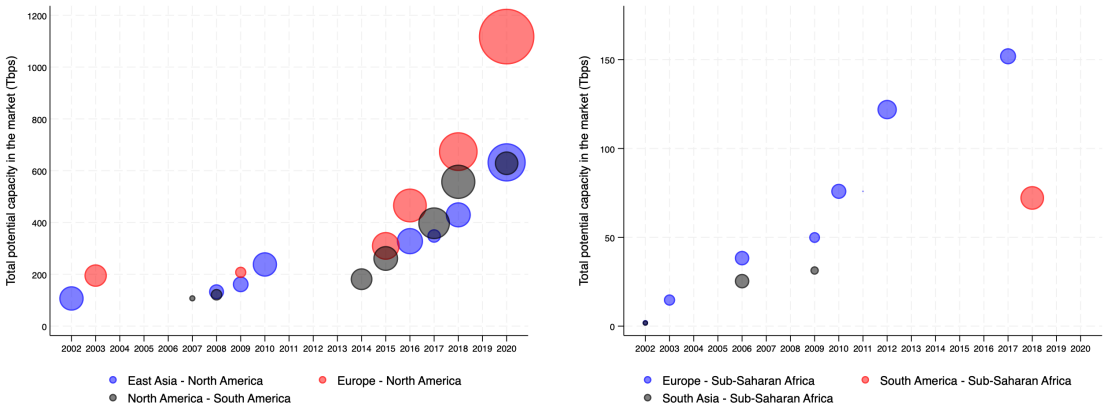
(d) Number of markets served by firm

Notes: Panel (a) shows the number of new cables and newly added capacity. Panel (b) shows the number of investments from 2002 to 2020 by firm conditional on investing at least once in this period. Panel (c) shows the number of cables owned by firm. Panel (d) shows the number of markets served by firm.

Figure 3: Investment episodes and total capacity in high and low capacity markets

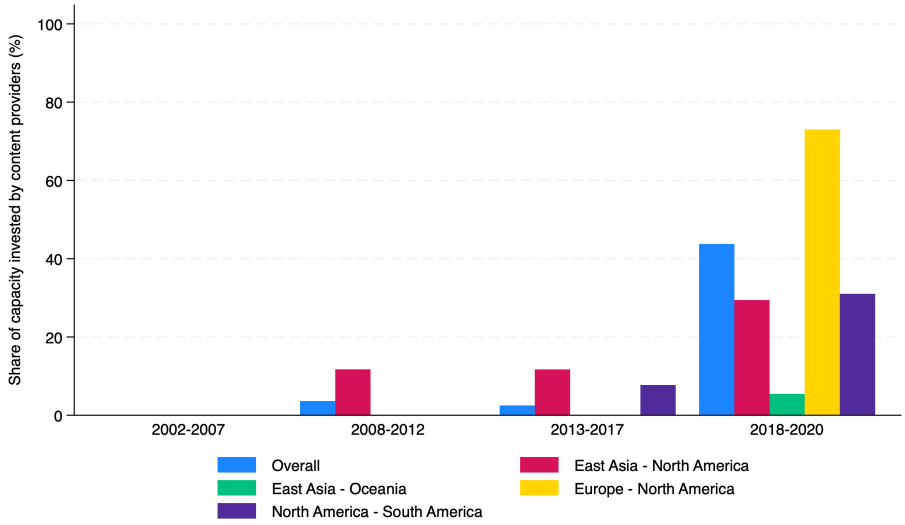
(a) High capacity markets

(b) Low capacity markets



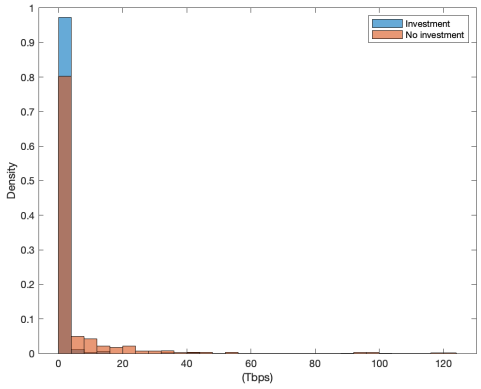
Notes: This figure shows the size of new investments represented by the size of the bubbles in relation to the total existing capacity for three high capacity markets and three low high capacity markets separately.

Figure 4: The share of newly added capacity by content providers by market

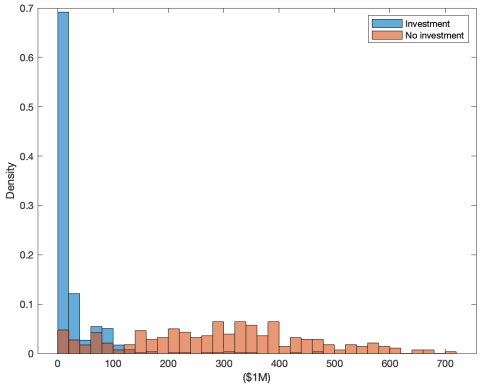


Notes: This figure shows the share of capacity invested by content providers for the overall industry and for the markets with the highest share of content provider investment in 2020.

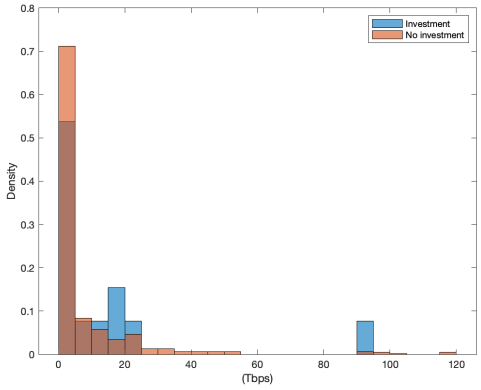
Figure 5: Data patterns helping identification



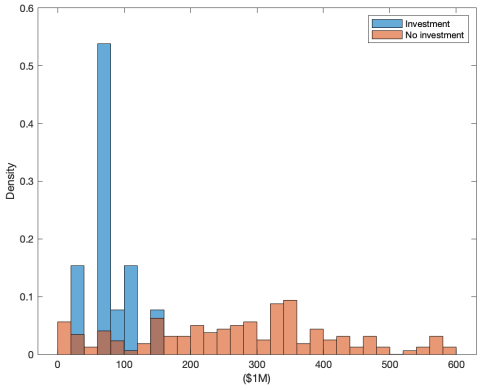
(a) Carriers: $\sum_m \Delta q_{jmt}$



(b) Carriers: Δc_{jt}



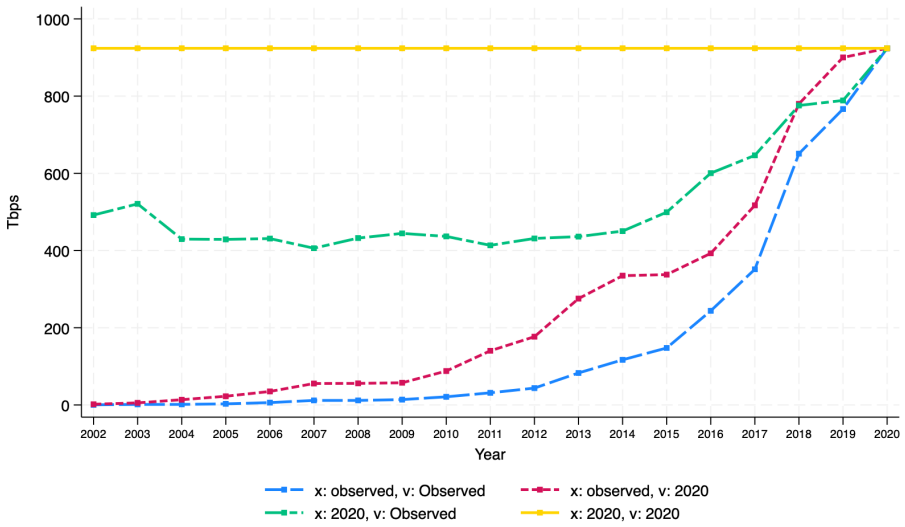
(c) Content providers: $\sum_m \Delta q_{jmt}$



(d) Content providers; Δc_{jt}

Notes:

Figure 6: The decomposition of trade in data



Notes: This figure plots the total data flows computed based on our model estimates, holding country-pair characteristics (x_{ckt}) and quality of connection (v_{ckt}) at the observed levels. It also plots data flows under the following three counterfactuals: (i) holding x_{ckt} and v_{ckt} at the 2020 levels; (ii) holding x_{ckt} at the 2020 level and v_{ckt} at the observed level; and (iii) holding v_{ckt} at the 2020 level and x_{ckt} at the observed level.

Table 1: Summary statistics

	Mean	SD	5th	95th	N
Cable attributes					
Number of owners	6.30	8.18	1	20	102
Number of markets served	1.30	0.82	1	3	102
Cable length (in km)	10,712.96	7,532.99	1,300	25,000	101
Potential capacity (in Tbps)	33.76	49.08	1	142.56	102
Construction costs (in million USD)	467.00	398.25	26	1,300	89
Direct cable	0.86	0.35	0	1	102
Owner attributes					
Number of markets served	4.15	3.82	1	12	2798
Number of cables owned	2.91	3.41	1	9	2798
Total potential capacity (in Tbps)	3.02	9.43	24	14.73	2798
Market attributes					
Total potential capacity (in Tbps)	73.10	133.20	0	397	399
Number of cables	3.64	4.55	0	14	399
Average length of direct cables (in thousand km)	10.10	6.12	2	21	161
Owner-level HHI of potential capacity	1,948.10	2,287.88	208	6,929	285
Number of owners	29.07	23.98	0	66	399

Table 2: International bandwidth usage

Year	(1) Total international bandwidth usage (Tbps)	(2) Number of countries making up 50% of usage	(3) Number of countries making up 90% of usage	(4) HHI
2002	2.97	4	14	1076.32
2003	4.83	3	14	1141.85
2004	6.83	4	15	1066.20
2005	10.00	4	17	1041.90
2006	14.00	4	20	978.47
2007	22.97	4	20	970.06
2008	38.45	4	25	860.42
2009	61.08	4	26	798.88
2010	93.37	5	28	777.18
2011	139.98	5	28	753.24
2012	202.10	5	29	739.92
2013	290.89	5	29	719.61
2014	424.16	5	28	720.43
2015	599.21	5	28	721.82
2016	888.17	5	28	718.06
2017	1,335.16	5	27	730.33
2018	1,995.80	5	26	736.50
2019	2,918.79	5	24	767.32
2020	4,149.75	5	24	719.16

Notes: Column (1) of this table shows the total international bandwidth usage from 2002 to 2020. Columns (2) and (3) show the number of countries that make up 50% and 90% of usage, respectively. Column (4) shows the HHI of usage at the country level.

Table 3: The effect of a new cable on usage growth

Dependent variable: % change in usage	(1) t to $t+1$	(2) t to $t+1$	(3) $t-1$ to $t+1$	(4) $t-1$ to $t+1$
Invest $(t-1)$	0.097***	0.094***	0.169***	0.170***
	(0.032)	(0.032)	(0.055)	(0.054)
Lagged number of cables $(t-2)$	0.021***		0.038***	
	(0.004)		(0.007)	
Invest x Lagged number of cables	-0.010***		-0.015**	
	(0.004)		(0.006)	
Lagged log total capacity $(t-2)$		0.070***		0.138***
		(0.015)		(0.025)
Invest x Lagged total capacity		-0.010***		-0.015**
		(0.003)		(0.006)
Observations	229	197	229	197
R-squared	0.768	0.778	0.805	0.821
Year FE	X	X	X	X
Market FE	X	X	X	X

Table 4: Path decomposition for Europe - North America

Market	Origin region	Destination region	Predicted Demand (Gbps)	Share (%)
Europe - North America	Europe	North America	209,185.91	83.95
Europe - North America	East Asia	North America	14,841.92	5.96
Europe - North America	North America	South Asia	6,675.32	2.68
Europe - North America	East Asia	South America	5,332.25	2.14
Europe - North America	North America	Taiwan	2,195.65	0.88
Europe - North America	North America	South Korea	1,956.92	0.79
Europe - North America	Japan	North America	1,650.14	0.66
Europe - North America	North America	Singapore	1,476.17	0.59
Europe - North America	North America	Philippines	1,220.20	0.49
Europe - North America	Europe	Japan	1,144.77	0.46

Table 5: Model estimates

	Estimates	(SE)
Panel A: Demand for data (x_{ckt})	θ^d	
Constant	-219.91	(8.28)
GDP	3.91	(0.49)
Population	-1.61	(0.47)
Time trend	0.54	(0.23)
2008-2012	-0.98	(0.84)
2013-2017	-0.26	(0.91)
2018-2020	0.50	(0.79)
Panel B: Cable usage (Z_{pckt})	θ^δ	
Constant	85.86	(9.78)
Path length	-3.52	(1.00)
Potential capacity	1.15	(0.35)
Number of cables	0.37	(0.12)
Panel C: Markup	γ ($m/ Tbps)	
	Non-content providers	Content providers
2002-2007	(5230.00 , 7760.00)	NA
2008-2012	(350.00 , 460.00)	(5.10 , 29.40)
2013-2017	(23.10 , 66.90)	(2.20 , 13.80)
2018-2020	(11.10 , 12.40)	(4.60 , 12.30)

Notes: Panels A and B of this table report estimates for θ^D, the parameters that govern how country-pair characteristics (x_{ckt}) affect demand for data transmissions between the two countries and estimates for θ^δ, the parameters that govern how cable features of a path (Z_{pckt}) affect the share of the data served by that path. The variable 'GDP' is computed as the sum of logged GDP for two countries in the country pair and the variable 'Population' is constructed similarly. The specification also includes a linear time trend, indicators for the 2008-2012, 2013-2017, and 2018-2020 periods, and market fixed effects. The 2002-2007 period is the omitted category. The path length is computed as the sum of cable lengths over markets involved in the path. We take the minimum capacity over markets as the potential capacity of a path and the capacity-weighted average as the number of cables. In Panel C, we report the 95% confidence set for the markup parameter, γ, capturing the variable profit (measured in millions of US dollars) from an increase in the data flows captured by a firm $\left(D_{mt}\frac{k_{jmt}}{K_{mt}}\right)$ by one Tbps.

Table 6: Counterfactuals: Adding a new cable in selected markets in 2015

Panel A: Based on the lower bound of γ

Market	(1)	(2)	(3)	(4)	(5)
	Change in incumbent revenue			Entrant revenue	Business stealing -(3)/(4)
	Investment market	Other markets	Total		
Overall	-80.87	13.31	-67.56	117.59	0.57
E. Asia - N. America	-94.40	-31.75	-126.16	75.92	1.66
Europe - N. America	-137.10	-28.40	-165.50	269.35	0.61
E. Asia - Sub-Saharan Africa	-0.21	6.21	6.00	17.94	-0.33
S. America - Sub-Saharan Africa	-0.10	11.73	11.63	6.11	-1.90

Panel B: Based on the upper bound of γ

Market	(1)	(2)	(3)	(4)	(5)
	Change in incumbent revenue			Entrant revenue	Business stealing -(3)/(4)
	Investment market	Other markets	Total		
Overall	-234.36	38.70	-195.66	340.57	0.57
E. Asia - N. America	-275.48	-91.97	-367.45	219.87	1.67
Europe - N. America	-397.06	-81.58	-478.64	780.06	0.61
E. Asia - Sub-Saharan Africa	-0.61	17.99	17.38	51.96	-0.33
S. America - Sub-Saharan Africa	-0.28	33.99	33.71	17.71	-1.90

Notes: The revenues are reported in millions of US dollars. Business stealing is computed as the ratio of the negative value of the change in the incumbent revenue in column (3) to the entrant revenue in column (4).

A Additional Figures and Tables

Table A1: Construction Cost Regression

Construction Cost	(1)	(2)
Log(Capacity)	0.0656	0.112**
	(-0.0483)	*(-0.0498)
Log(Cable Length)	1.096***	1.075***
	(-0.0668)	(-0.0503)
Time trend		-0.0292**
		(-0.0118)
Constant	9.060***	9.086***
	(-0.74)	(-0.646)
Observations	31	31
R-squared	0.887	0.908

Table A2: Capacity Regression

	Log Capacity
Year fixed effects	Yes
Market fixed effects	Yes
Constant	8.862***
	-0.82
Observations	81
R-squared	0.684

B Inference Procedure

We implement the conditional chi-squared (CC) test propsed by Cox and Shi (2023) to compute confidence sets for the true parameter value μ^* based on a set of unconditional moment inequalities $l = 1, ..., L$. We denote the set of L moments we use

in the estimation as

$$\overline{m}(\gamma) = (m_1(\gamma), ..., m_L(\gamma))'$$

with the inequalities given by

$$\overline{m}_l(\gamma) \geq 0, l = 1, ..., L.$$

$\hat{\Sigma}(\theta)$ is an estimator of variance-covariance matrix of the moments. The procedure looks as follows.
Step 1: Define a grid $\mathcal{A}$ that will contain the confidence set.
Step 2: Choose a point in the grid $\gamma \in \mathcal{G}$. For a given significance level $\alpha \in (0, 1)$, the following steps test the null hypothesis given by $H_0 : \gamma^* = \gamma$.
Step 3: Evaluate the quasi-likelihood ratio statistic at γ as

$$T(\gamma) = \min_{\mu:\mu\geq 0} N(\overline{m}(\gamma) - \mu)'\hat{\Sigma}(\gamma)^{-1}(\overline{m}(\gamma) - \mu) \tag{11}$$

where N is the sample size. The vector μ is of dimensions $L1$. When defining the objective function of the minimization problem in equation (11), it occasionally happens that the covariance matrix $\hat{\Sigma}(\theta_p)$ is singular or close to singular. When this is the case, we follow Andrews and Barwick (2012), and substitute the covariance matrix in equation (11) for the following matrix:

$$\tilde{\Sigma}(\gamma) = \hat{\Sigma}(\gamma) + \max\{0.012 - \text{det}(\hat{\Omega}(\gamma), 0\} Diag(\hat{\Sigma}(\gamma) \tag{12}$$

where $Diag(\hat{\Sigma}(\gamma))$ is the $L \times L$ diagonal matrix whose diagonal elements are equal to those of $\hat{\Sigma}(\gamma)$ and $\hat{\Omega}(\gamma)$ is the correlation matrix of the moments evaluated at γ.

Step 4: Count how many values of μ equal 0. We denote this number as $\hat{r}$.
Step 5: Accept/reject θ_p. Include γ in the $(1-\alpha)\%$ confidence set, $\hat{G}^{1-\alpha}$, if $T(\gamma) \leq \chi^2_{\hat{r},1-\alpha}$, where $\chi^2_{\hat{r},1-\alpha}$ is the $100(1-\alpha)\%$ quantile of $\chi^2_{\hat{r}}$, the chi-squared distribution with $\hat{r}$ degrees of freedom.
Step 6: Repeat steps 2 to 5 for every γ in the grid