Investment and Usage of the Subsea Internet Cable Network*

Jihye Jeon[†] and Marc Rysman[‡]

April 2025

Abstract

This paper studies the construction and use of undersea internet cables. In our model, firms invest in new cables that alter global data flows. We model countries exchanging data and how these data flows traverse the cable network. We estimate this model using new data on cable construction and usage via moment inequalities. We show that improvements in the cable network are an important contributor to growth in internet usage relative to secular trends. Our counterfactuals show the potential for an inefficient allocation of cables and highlight network spillovers as more important than business stealing in driving this result.

*We thank Haoran Pan, Tim Lee, and Max Levine 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. We thank Giulia Brancaccio, Nicholas Vreugdenhil, and the seminar and conference participants at the ASSA Annual Meeting, Boston University, Columbia University, EUI, Federal Communications Commission, Harvard Business School, IIOC, the Luiss Business School, NBER IO Meeting, Universidad Autónoma de Barcelona, New York University, North Carolina University, Princeton University, Telecom Policy Research Conference, the University of Georgia, the University of Toronto, the Conference on Auctions, Firm Behavior, and Policy at the University of Oklahoma, and the Alumni Conference at the University of Wisconsin.

[†]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 overseas 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 ask how much investment in the cable network has contributed to the growth of internet usage. We decompose the dynamics of internet usage into the amounts due to cable network growth and to secular and demographic trends in internet usage. This research question highlights the importance of subsea cable investment.

Second, we ask if investment in the cable network is allocated efficiently. The subsea internet cable market is almost exclusively owned by private investors and is subject to a very low level of economic regulation. Thus, there are two primary reasons that the observed network may exhibit inefficiency. One is classic business stealing - new investment focuses on large markets that typically already have high investment because new investment will steal market share from incumbents. The second reason is the networked nature of the product: investment in one market may stimulate demand in connected markets but the investor will not naturally capture this benefit, leading to externalities across the network. As a result, the network we observe is unlikely to exhibit global efficiency. We develop a model that highlights both the international implications of local investments and the role of inefficient business stealing incentives in investment decisions.

Our data is primarily drawn from TeleGeography, a proprietary data company that covers internet equipment and services, with particular expertise in the subsea 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. 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 data set goes from 2002 to 2020. We supplement the data with various sources including trade data from CEPII. We develop several motivating facts from our data set. For instance, even though data flows become less geographically concentrated over time, the concentration of cable capacity remains constant, suggesting that investors overemphasize large markets.

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 interregional used bandwidth. 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 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, busy routes are often central to global communication, so the social value of investment may be high because of demand enhancement in adjacent and distant markets. How these forces offset each other determines whether the private allocation of investment is optimal.

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. Both demographics and measures of the quality of the cable connection affect the level of data flows between any two countries. The estimates suggest that an increase in the sum of GDP per capita of the involved countries would lead to an increase in data flows, and there is a strong positive time trend. 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 utilize moment inequalities to estimate profits that cable owners collect on data flows. 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 content and non-content providers (referred to as *carriers* in the rest of the paper) to study how their incentives differ.

We recover a 95% confidence interval for the markup for three distinct time periods separately and for content providers and carriers. We find that markups fall rapidly in our sample period. For example, we estimate that a one-Tbps increase in bandwidth served by a carrier would result in a profit increase between 2.7 to 3.5 billion dollars in 2002-2009, and 12.0 to 15.6 million dollars in 2018-2020. This is consistent with the steep decline in bandwidth pricing observed during this period. We find that content providers' markups are comparable to carriers'.

Taking the estimates together, we conduct two main counterfactual analyses. 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 data flows, on par with increases in demand. In 2013, for example, improving the quality of the cable network to the 2020 level while fixing the demographics at the observed level would increase the total data flows by 6 times from the observed level, or bring it up to 26% of the 2020 level.

In our second set of counterfactuals, we characterize the socially optimal cable network. We find that a planner who takes into account externalities across markets as well as competition externalities would add a new cable to 8 out of 15 markets we consider at the end of 2020. These include not only markets that have large demand and high existing cable capacity, such as Europe-North America, but also small markets with little existing capacity, such as Europe-South America and East Asia-Sub-Saharan Africa. To understand the extent and sources of inefficiencies, we consider a de novo entrant constructing a new cable to each market. Comparing entrant profits to changes in total surplus in each market, we find that although these two measures are highly correlated, there are important cases in which the two are misaligned. In markets like Europe-North America, a new cable generates large market expansion across many markets, producing a large social surplus gain, while also guaranteeing a high profit to an entrant. However, in East Asia-Sub-Saharan Africa, for example, although a new cable would generate substantial social value, entrants would not find it optimal to invest in it.

Even though business stealing externalities are substantial in many markets, since they are strongest in large markets where a new cable is socially valuable, they are not critical in creating mismatches between private and social incentives. We instead find that network externalities play a more critical role. When a new cable arrives in markets like East Asia-Sub-Saharan Africa, a high share of market expansion happens outside the market. Thus, it does not create enough revenue to the entrant. We find that even incumbents who own a large cable network, thus internalizing some of the out-of-market externalities, would not find it profitable to invest in this market. Because of the global nature of the externality, even individual national governments may not be sufficiently incentivized to address the difference between social and private gains.

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. Cariolle (2021) is similar, focusing on the local internet access after the arrival of a cable in Africa. Eichengreen, Lafarguette, Mehl and Minesso (2023) show that the arrival of a new cable in Singapore affected currency trading. The closest paper that we are aware of is the concurrent paper by Caoui and Steck (2025), which focuses on the role of path diversity. Whereas their paper models dynamic decision-making, our model takes into account the fact that all regions are globally connected, which means new construction in one market can affect data flows in others. In separate research, we characterize data flows on subsea cables as *international trade in data* and we contrast the evolution of trade in data with trade in goods (Jeon and Rysman, 2024). Cariolle, Houngbonon, Silue and Strusani (2024) study how cable connectivity impacts downstream broadband prices. More broadly, our paper relates to papers that study infrastructure that supports the internet such as Greenstein (2015).

There is also a literature outside of economics that studies subsea internet cables. Particularly relevant for us is Fanou, Huffaker, Mok and Claffy (2020), which uses internet measurement techniques to find significant benefits in terms of reduced latency arising from the first major cable investment between South America and Africa in 2018.

We find that distance matters for internet communication, contributing to a broad literature with similar findings. In addition to Fanou et al. (2020) already mentioned, Smith (2009), Leighton (2009), and Clark and Wedeman (2024) are examples in the computer science literature that discuss the importance of distance in determining internet performance. In the economics literature, LaRiviere, Kannan and Wang (2020) show that cloud computing consumers value nearby computing services. Elliott, Houngbonon, Ivaldi and Scott (2024) and Malone, Nevo and Williams (2021) find consumer responsiveness to investment in internet facilities.

Our emphasis on network infrastructure is a key feature 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). Examples of equilibrium models of transportation and cargo networks are Fréchette, Lizzeri and Salz (2019) and Brancaccio, Kalouptsidi and Papgeorgiou (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, Newberry and Seim (2022), and in airlines, such as Aguirregabiria and Ho (2010). 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).

Lastly, we contribute to a growing empirical literature that employs moment inequalities, recently reviewed by Molinari (2020), Kline, Pakes and Tamer (2021) and Canay, Illanes and Velez (2023). Our paper is especially closely related to methods in Pakes (2010) and Pakes, Porter, Ho and Ishii (2015) and 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 (2024).

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 have 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 is 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. For example, there is no regulation of pricing or interconnection. There is some government intervention in regard to national security, but that has little impact on the features of the market we consider.[1] 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 used bandwidth. 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 papers by Caoui and Steck (2025) and Cariolle et al. (2024). 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 terrestrial cables. Figure 1 provides a map from TeleGeography of commercial subsea cables in 2023.

[1]For example, in the United States, all subsea cable landings must be approved by the Federal Communications 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, which became a formal committee in 2020. 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.

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

Figure 1: Map of subsea cables in 2023

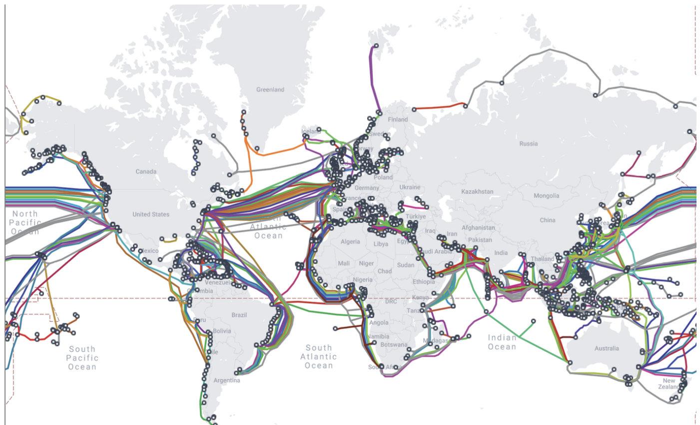

Notes: This map shows undersea telecommunications cables as of October 5, 2023. Source: TeleGeography, Submarine Cable Map, www.submarinecablemap.com

We observe several measures of capacity at the cable level. *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*, or equivalently, *used bandwidth*. TeleGeography constructs an equivalent to used bandwidth 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 bandwidth as our measure of usage (as do the other papers working with these data). In general, purchasers of capacity devote substantial engineering resources to tightly manage the relationship between purchased capacity, used bandwidth, traffic, and realized quality of communication, so used bandwidth is likely a good approximation of traffic.

[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.

TeleGeography offers used bandwidth price data at the city-pair level for select cities. We do not use these data in our main analysis in part because the pricing data cover only a small subset of city-to-city pairs, which would limit our ability to study the global market.[4] However, we use the pricing data to calculate some "back-of-the-envelope" measures of consumer surplus in our counterfactuals section.

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.[5]

We divide the world into the following seven regions: East Asia, Europe, North America, Oceania, South America, South Asia, and Sub-Saharan Africa. TeleGeography 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 Himalayan 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.

We define *markets* as pairs of regions (e.g. North America-Europe, North America-South America). Whereas TeleGeography reports capacity at the level of the cable, it reports used bandwidth annually at the level of the region pair.[6] We observe

[4]Other issues for our purposes are that 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. Traffic between two cities that do not have those cities as endpoints may not pay this price. To the extent traffic between cities uses terrestrial networks (which must be the case if cities are inland or do not have a cable landing point), city-to-city prices reflect those costs as well.

[5]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).

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

a single used bandwidth number for the region pair, the total of used bandwidth going in both directions.[7] We observe used bandwidth between each of the seven regions. However, if there is no cable between two regions, used bandwidth for that region pair is naturally zero. Importantly, used bandwidth between two regions may reflect traffic that is from further endpoints. For example, there was no major cable between Europe and South America as of 2020.[8] 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. We observe the used bandwidth between any two region pairs but not the endpoints of the associated traffic. Throughout the paper, we refer to *used bandwidth* at the region-pair level as a measure of cable usage through the two regions, including traffic between endpoints besides the two regions. In contrast, we use *data flows* to refer to endpoint-to-endpoint usage between two countries.

TeleGeography also reports international bandwidth usage at the country level (not at the country pair level), that is, the country's total used bandwidth to other countries. 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. One issue with this variable is that it includes usage on terrestrial cables, which makes it difficult to integrate into our model. To overcome this issue, we leverage the fact that, for island countries, all international usage must be on subsea cables. We identify seven island countries with substantial used bandwidth: Singapore, Japan, Taiwan, Australia, South Korea, the Philippines, and New Zealand.[9] Section 4.2 describes how including islands as separate regions provides additional moments in our estimation.[10]

[7] 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.

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

[9] 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.

[10] 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,

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 interregional cables, which connect two regions, or a region and an island, or two islands. For instance, we drop cables that connect Italy to Greece or connect among Indonesian islands.

We supplement TeleGeography's data with gravity data from Centre d'Etudes Prospectives et d'Informations Internationales (hereafter CEPII). This data set provides country-level and country-pair-level information 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: Summary statistics

	Mean	SD	5th	95th	N
Panel A. 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	0	143	102
Construction costs (in million USD)	467.00	398.25	26	1,300	89
Panel B. Owner attributes					
Number of markets served	4.15	3.82	1	12	2798
Number of cables owned	2.91	3.41	1	9	2798
Number of investments	0.18	0.67	0	1	2798
Total potential capacity (in Tbps)	3.02	9.43	0	15	2798
Panel C. Market attributes					
Total potential capacity (in Tbps)	73.10	133.20	0	397	399
Number of cables	3.64	4.55	0	14	399
Number of owners	29.07	23.98	0	66	399

Notes: Panel A summarizes cables in our sample–interregional cables. The observations are at the cable level. Panels B and C summarize cable owners and markets (pairs of regions) where the observations are at the owner-year and market-year levels, respectively. The data include 102 unique cables, 197 owners, and 21 markets.

Table 1 provides information on the final sample of the cables, owners, and markets

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 North America and Europe (potentially with other endpoints) but not how much goes between the US and France.

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. Our sample includes 197 unique owners. On average, an owner owns 2.9 cables and 3 Tbps of capacity and is present in 4.1 markets. Our sample has 21 unique markets of which 15 have at least one cable. 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.

As shown in Figure A.1(a), there is active investment throughout the sample period, with a slight rise in the number of investment episodes in the 2016-2020 period. The total capacity added in this period is much larger as well, partially due to the increase in the average capacity of new cables. Figure A.1(b) shows that used bandwidth grew dramatically from 2002 to 2020 by a factor of 1000.

2.3 Motivating Facts

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 2, we plot total capacity and number of investment episodes by market.[11] We find that new investments tend to be concentrated in certain markets 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, while there is limited investment in markets such as Europe-South America, South America-Africa, and East Asia-Africa.

Figure 3(a) provides another way to look at this issue. The figure computes the HHIs of used bandwidth and capacity across markets by year. For example, the HHI for used bandwidth will be high if used bandwidth is concentrated in a few markets. Comparing the two lines, we see that even though used bandwidth became substantially less geographically concentrated from 2002 to 2020, the concentration of cable capacity remained stable, suggesting possible overinvestment in large markets.

[11] We omit six markets that have no active cable in our sample period from this plot. They include North America-Sub-Saharan Africa, North America - South Asia, Oceania - South America, South America - South Asia, Oceania - Sub-Saharan Africa, and East Asia - South America.

Figure 2: Total capacity and number of newly added cables by market

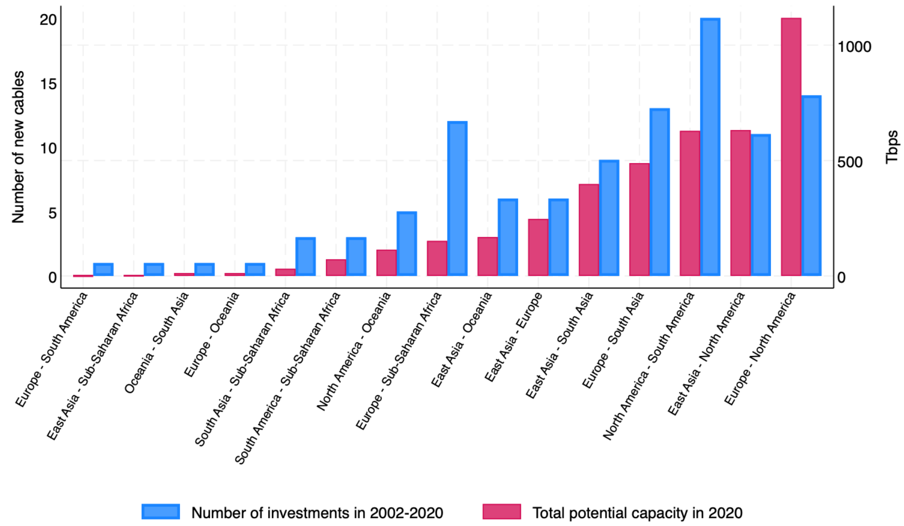

Notes: This figure plots total potential capacity in 2020 and total number of newly added cables from 2002 to 2020 by market. The bars are ordered by total capacity.

Various factors may contribute to these patterns of investment, including heterogeneity in demand, business stealing incentives, and geographical spillovers. Our model proposed in Section 3 allows us to distinguish these factors.

Figure 3: Data patterns on cable investment and used bandwidth

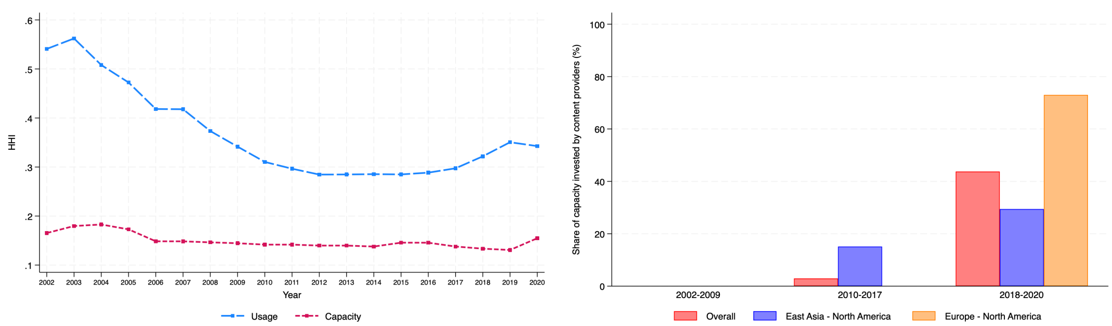

(a) HHI of usage and capacity (b) Share of investment by content providers

Notes: Panel (a) shows the HHI of market-level used bandwidth and cable capacity. To compute the HHI, we compute the share of used bandwidth (capacity) by dividing each market's used bandwidth (capacity) by the total used bandwidth (capacity) in all markets. Then, we sum the squared shares across all markets. Panel (b) shows the share of capacity invested by content providers for the overall industry and for the East Asia-North America and Europe-North America markets (that have the highest share of content provider investment in 2020).

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 co-investors, 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 (now Meta), Amazon, and Microsoft. Figure 3(b) shows the share of capacity invested by these content providers for the three 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 remained very small (below 3% in 2010-2017). 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.

The investment patterns guide our estimation in two ways as described in more detail in Section 4. First, we allow the profit function to be different for content providers. Second, we group our sample period into three periods: 2002-2009, 2010-2017, and 2018-2020. This coincides with overall capacity increases as seen in Figure A.1(a). There are relatively low levels of capacity added from 2002-2009, an uptick from 2010 to 2017, and very high levels in 2018-2020.[12] In addition, 2010 marks the first investment by content providers and the rate of investment by content providers is distinctly higher in 2018-2020. Thus, we include period fixed effects in the demand for data flows and allow the profit parameters to be different across these three periods.

3 A Model of Data Flows and Cable Investment

3.1 Cable Investment

In each period t, firms $j = 1, \ldots, J$ simultaneously choose whether to invest in a new cable in each market $m = 1, \ldots, M$. Markets are non-directional pairs of geographic

[12]The low level of investment starting in 2002 was known in the US as a result of the dot.com bubble of 2000, which was associated with a glut of dark fiber.

regions, and investment connects regions via an undersea cable. Firms profit from their ownership share of cables and the amount of data that traverses those cables. We model firms as myopic and discuss robustness checks on this issue below.

In each period, a new line may be constructed in a market. Each firm simultaneously makes a binary choice $a_{jmt} \in \{0, 1\}$ in each market whether to build the available line. Let $\boldsymbol{a}_{mt} = (a_{jmt})_j$ denote the vector of choices in market m and let $\boldsymbol{A_t}$ be the $J \times M$ matrix of choices in all markets. The capacity of the potential cable is a known exogenous function of investment choices $\tilde{K}_{mt}(\boldsymbol{a}_{mt})$. For instance, if $a_{jmt} = 0$ for all j, then $\tilde{K}_{mt}(\boldsymbol{a}_{mt}) = 0$. That is, if no firms invest, no new capacity is added to the market. The function $\tilde{K}_{mt}(\cdot)$ is subscripted by mt to allow the capacity of new cables to vary across markets and time periods, for instance, growing in later periods.

If multiple firms choose to invest in a given market, they become joint owners of the line, evenly splitting ownership of the line. Let $\tilde{k}_{jmt}(\boldsymbol{a}_{mt})$ be the amount of new capacity that firm j owns in market mt: $\tilde{k}_{jmt}(\boldsymbol{a}_{mt}) = \frac{a_{jmt}}{\sum_{z=1}^{J} a_{zmt}} \tilde{K}_{mt}(\boldsymbol{a}_{mt})$. Total capacity in market j is:[13]

$$K_{mt}(\boldsymbol{a}_{mt}, K_{mt-1}) = K_{mt-1} + \tilde{K}_{mt}(\boldsymbol{a}_{mt}), \tag{1}$$

and firm j's capacity is:

$$k_{jmt}(\boldsymbol{a}_{mt}, k_{jmt-1}) = k_{jmt-1} + \tilde{k}_{jmt}(\boldsymbol{a}_{mt}). \tag{2}$$

We denote the total cost of the line as $C_{mt}(\boldsymbol{a}_{mt})$, where total cost depends on choices through $\tilde{K}_{mt}(\boldsymbol{a}_{mt})$. Investors evenly split the cost of the line. The cost to firm j is:

$$c_{jmt}(\boldsymbol{a}_{mt}) = \frac{a_{jmt}}{\sum_{z=1}^{J} a_{zmt}} C_{mt}(\boldsymbol{a}_{mt}). \tag{3}$$

We also track the number of cables constructed in a market. Let n_{mt-1} be the number of cables in market m in the previous period. The process for this variable is:

$$n_{mt}(\boldsymbol{a}_{mt}, n_{mt-1}) = n_{mt-1} + \mathbb{1}\{a_{jmt} = 1 \text{ for any } j = 1, \ldots, J\}. \tag{4}$$

[13] We have written capacity accumulation as if there is no depreciation. In practice, some cables are eventually decommissioned and are then no longer used. We account for this in our data construction. Cables typically last more than 20 years, at which point their capacity is very small relative to what is in use so accounting for decommissioning has a minimal effect on our empirics.

Let the set of state variables $\mathcal{S}_t = \{\{k_{jmt-1}\}_{j=1,\ldots,J,\, m=1,\ldots,M}, \{n_{mt-1}\}_{m=1,\ldots,M}\}$. Note that $\mathcal{S}_t$ naturally captures K_{mt-1} for all markets. The profit for firm j is given by:

$$\bar{\pi}_{jt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right) = \sum_{m \in \mathcal{M}} \left\{ r_{jmt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right) - c_{jmt}\left(\boldsymbol{a}_{mt}\right) + \nu_{ajmt} \right\},$$

where $r_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t)$ is the revenue to firm j in market mt. Revenue here is net of variable costs, i.e. revenue is total profit before paying the investment cost. As will be clear in our specification of demand, revenue in each market depends on capacity and the number of cables in each market, which depend on global actions $\boldsymbol{A}_t$ and states $\mathcal{S}_t$. The term ν_{ajmt} is a shock to the firm profit observed by the firm, but unobserved by the econometrician. In our framework, which is based on Pakes et al. (2015), the value of ν_{ajmt} can be common knowledge across firms or privately observed.

Firms have measurement error over elements of $r_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t)$, 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 $\{a_{jmt}\}_m$, firm j perceives the profit from investment in mt as:

$$\begin{aligned} \pi_{jt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right) = &\ E\left[\bar{\pi}_{jt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right) \middle| \mathcal{J}_{jt}\right] \\ = & \sum_{m \in \mathcal{M}} \left\{ r_{jmt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right) + \varepsilon_{ajmt} - c_{jmt}\left(\boldsymbol{a}_{mt}\right) + \nu_{ajmt} \right\}. \end{aligned} \tag{5}$$

Firms play a Nash equilibrium, simultaneously solving: $\max_{a_{jmt},\, m=1,\ldots,M} \pi_{jt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right)$. Recall that a_{jmt} is element jm of $\boldsymbol{A}_t$.

We specify the revenue function as:

$$r_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t) = \gamma_{jt} q_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t) = \gamma_{jt} D_{mt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right) \frac{k_{jmt}\left(\boldsymbol{a}_{mt}, k_{jmt-1}\right)}{K_{mt}\left(\boldsymbol{a}_{mt}, K_{mt-1}\right)}, \tag{6}$$

where γ_{jt} is the parameter of our interest and $q_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t)$ denotes the quantity served by firm j in mt. We denote the total data transmitted in market m as $D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t)$. It depends on capacity in all markets.

The cost function is drawn from observable cost data, as described below in Section 4.3. We specify the data function $D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t)$ below in Section 3.2.

3.2 Data Flows and Cable Usage

In this section, we present an estimable model of how country-to-country data flows maps into used bandwidth between regions that we observe as used bandwidth. Countries are indexed by $c = 1, \ldots, C$ and each country pair ck has data flows between each other. Data flows 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 data flows transmitted between countries are a function of their demographic and country characteristics, and the quality of the connection between the two countries, which is a function of cable characteristics. The total data flows between the two countries is:

$$d_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \frac{\exp(x_{ckt}\theta^d + v_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t))}{1 + \exp\left(x_{ckt}\theta^d + v_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t)\right)} \overline{M}, \tag{7}$$

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 $v_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ 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 *regions* $r = 1, \ldots, R$. The regions form a partition of the countries. As described above, regions are chosen so that communication between countries in different regions must traverse subsea cables. Recall that we denote pairs of regions as *markets* indexed by $m = 1, \ldots, M$.

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 markets for which there is an active cable. 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 region 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.

We focus on paths that connect four regions at most. 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 and no intermediate points as the *direct path*. So, in the previous list, the first element is the direct path and the rest are indirect paths.

Data are allocated to paths by internet routers. In practice, this is a complex process for which our model provides a simple reduced-form approximation.[14] Each path is characterized by an *attractiveness* $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$. Attractiveness $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ depends on cable features such as path distance and capacity associated with each path as well as the number of cables on the path, which may reflect competition or the resiliency of the path. [15] The share of data going from c to k on path p at time t is:

$$s_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \frac{\exp\left(\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)\right)}{\sum_{z=1}^{P_{ckt}} \exp\left(\delta_{zckt}(\boldsymbol{A}_t, \mathcal{S}_t)\right)} \tag{8}$$

Attractiveness $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ depends on the level of capacity on the path, which we denote as $K^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$. Here, superscript "$\delta$" indicates a path-level variable. Naturally, $K^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ is a function of the capacity connecting the regions on path p. We assume $K^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ is the minimum of these capacities, which captures that the lowest capacity leg creates a bottleneck. Formally, let $\mathcal{M}_{pckt}$ be the set of markets on path p between country pair ck. We assume $K^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \min_{m \in \mathcal{M}_{pckt}} K_{mt}(\boldsymbol{a}_{mt}, K_{mt-1})$, where $K_{mt}(\boldsymbol{a}_{mt}, K_{mt-1})$ is defined in Equation (1). We have experimented with other specifications, such as a geometric average, but this specification seems most realistic. Appendix B provides further details on how we construct path capacity.

Similarly, $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ depends on the number of cables on the path $n^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$. We define $n^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ to be the capacity-weighted average of the number of cables on each market on the path. We parameterize the attractiveness of the path $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ to be a linear function of path characteristics as follows:

$$\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \theta^{\delta,0} + \theta^{\delta,Z} Z_{pckt} + \theta^{\delta,K} K^{\delta}_{pckt}(\boldsymbol{A}_{mt}, \mathcal{S}_t) + \theta^{\delta,n} n^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t),$$

[14] In general, internet routers automatically direct packets of information to the least crowded routes they can access. Even packets from the same overall stream, for instance, a video or email message, may take very different routes to their destination. Furthermore, whether a router can access a particular path can depend on negotiations between network owners.

[15] 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.

where Z_{pckt} is the path length. The parameters $\theta^{\delta} = \{\theta^{\delta,0}, \theta^{\delta,Z}, \theta^{\delta,K}, \theta^{\delta,n}\}$ are to be estimated.

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

$$v_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \ln\left(\sum_{p=1}^{P_{ckt}} \exp\left(\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)\right)\right).$$

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

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

In our model, more paths between two countries (higher P_{ckt}) and the increased quality of those paths (higher $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$, 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 short periods of time and we have annual data. However, the model captures congestion in the sense that better connectivity leads to more network use.

Summing $\hat{d}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ over all paths that include m provides the total used bandwidth between countries c and k traveling in market m in period t:

$$\tilde{d}_{mckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \sum_{p=1}^{P_{ckt}} \hat{d}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t) \mathbb{1}\{m \in \mathcal{M}_{pckt}\}.$$

Thus, total used bandwidth in market m in period t is:

$$D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t) = \sum_{c=1}^{C-1} \sum_{k=c+1}^{C} \tilde{d}_{mckt}(\boldsymbol{A}_t, \mathcal{S}_t). \tag{9}$$

In estimation, we match the predicted used bandwidth $D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t)$ to the observed used bandwidth in the data. In addition to being the dependent variable in our demand estimation, used bandwidth $D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t)$ enters profit in Equation (6).

4 Estimation

We discuss how we make use of island countries and then discuss the estimation of the demand for data flows. We then turn to estimating the investment game.

4.1 Island Countries

As described above, TeleGeography reports total international used bandwidth by country. For most countries, our demand model cannot be applied to this variable because we do not model the terrestrial network. However, for island countries, all used bandwidth must go over subsea cables so we can make predictions about country-level used bandwidth. We pick countries with substantial used bandwidth that do not have terrestrial connections to other countries, arriving at seven island countries. We separate regions with island countries into subregions such that there is one subregion for each island and the rest of the countries in the region become a separate subregion. We construct our data on total capacity, the number of cables, and other cable-level variables for each sub-region pair in the world rather than for each region. We can then solve for demand just as in Section 3.2, treating sub-regions as the unit of analysis rather than regions.

For this purpose, we let $\mathcal{W}_r$ be the set of markets containing region r and define total used bandwidth for a region as: $Q_{rt}(\boldsymbol{A}_t, \mathcal{S}_t) = \sum_{m \in \mathcal{W}_r} D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t)$. When we form moments using $Q_{rt}(\boldsymbol{A}_t, \mathcal{S}_t)$, we use only observations from our seven island countries. Since used bandwidth in our data is between regions, not between subregions, we still generate $D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t)$ at the region-pair level.

For the investment model, we use the region-level of analysis. That is, we add up subregions into regions and model investment between regions. We solve demand at the region level in order to compute predicted and counterfactual demand for the moment inequalities, further described below. We provide a mathematical restatement of the model accounting for subregions in Appendix C.

4.2 Data Flows and Used Bandwidth

On the demand side, we estimate parameters that determine country-to-country demand for data flows and the allocation of the demand across cable paths. The procedure involves matching the predicted used bandwidth to the observed levels in the data based on the generalized method of moments (GMM).

For this section, we write $D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t)$ as $D_{mt}(\theta)$ to emphasize the dependence on estimation parameters, where $\theta = \{\theta^d, \theta^\delta\}$. We analogously write $Q_{rt}(\theta)$. Let D^*_{mt} and Q^*_{rt} be the observed value in our data. Thus:

$$\xi^D_{mt}(\theta) = D^*_{mt} - D_{mt}(\theta) \qquad \xi^Q_{rt}(\theta) = Q^*_{rt} - Q_{rt}(\theta)$$
$$\Delta\xi^D_{mt}(\theta) = \xi^D_{mt}(\theta) - \xi^D_{mt-1}(\theta) \qquad \Delta\xi^Q_{rt}(\theta) = \xi^Q_{rt}(\theta) - \xi^Q_{rt-1}(\theta).$$

Let $\boldsymbol{y}^D_{mt}$ be a vector of instrumental variables. We assume that at the true parameters θ^0, $E[\xi^D_{mt}(\theta^0)|\boldsymbol{y}^D_{mt}] = 0$. Analogously, we assume $E[\xi^Q_{rt}(\theta^0)|\boldsymbol{y}^Q_{rt}] = E[\Delta\xi^D_{mt}(\theta^0)|\boldsymbol{y}^{D,\Delta}_{mt}] = E[\Delta\xi^Q_{rt}(\theta^0)|\boldsymbol{y}^{Q,\Delta}_{rt}] = 0$. Let $\boldsymbol{\xi}^D(\theta)$ be the $M \times T$ vector of elements $\xi^D_{mt}(\theta)$, and analogously for $\boldsymbol{\xi}^Q(\theta)$, $\boldsymbol{\xi}^{D,\Delta}(\theta)$ and $\boldsymbol{\xi}^{Q,\Delta}(\theta)$. Similarly, let $\boldsymbol{Y}^D$, $\boldsymbol{Y}^Q$, $\boldsymbol{Y}^{D,\Delta}$, and $\boldsymbol{Y}^{Q,D}$ be the matrices of instruments $\boldsymbol{y}^D_{mt}$ and so on. Then we have moments:

$$\boldsymbol{m}^D(\theta) = \boldsymbol{Y}^{D\prime}\boldsymbol{\xi}^D(\theta) \qquad \boldsymbol{m}^Q(\theta) = \boldsymbol{Y}^{Q\prime}\boldsymbol{\xi}^Q(\theta)$$
$$\boldsymbol{m}^{D,\Delta}(\theta) = \boldsymbol{Y}^{D,\Delta\prime}\boldsymbol{\xi}^{D,\Delta}(\theta) \qquad \boldsymbol{m}^{Q,\Delta}(\theta) = \boldsymbol{Y}^{Q,\Delta\prime}\boldsymbol{\xi}^{Q,\Delta}(\theta). \tag{10}$$

Recall that $\boldsymbol{Y}^Q$ and $\boldsymbol{Y}^{Q,\Delta}$ sum over sub-regions that contain a single island and drop other sub-regions. Let $\boldsymbol{\Gamma}(\theta)$ be the vector of stacked moments from Equation (10). The GMM estimator is then given by: $\hat{\theta} = \arg\min_\theta \boldsymbol{\Gamma}(\theta)'\mathbf{W}\boldsymbol{\Gamma}(\theta)$, where $\mathbf{W}$ is a positive definite weighting matrix that we select through two-stage GMM estimation.

We now discuss variation in data that allows us to estimate these parameters, including our choice of instruments. Generally, identification is challenging because we wish to infer both country-to-country demand and cable quality from observed used bandwidth. Taking the simplest example possible of two regions connected by a single path for one period, we would observe only one value of used bandwidth (that is, only one observation of the dependent variable). Naturally, we cannot separately infer both region-to-region demand and path quality in this case. More generally, even with an increased number of regions or time periods, a fully non-parametric treatment (that is, allowing $x_{ckt}\theta^d$ to freely vary across all combinations of region pairs and time as well as allowing δ_{pckt} to freely vary across all paths) 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 specify that country-to-country demand depends on demographic variables (e.g. GDP and population of the two involved countries). 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 be totally disconnected from the demand with another region. Moments drawn from island countries contribute to precision over country-level variables. Cable quality parameters in $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ are determined by the level of used bandwidth observed in markets conditional on country-to-country demand.

With regard to instrument choice, we face the same endogeneity issue as in many models of entry. That is, if firms can predict where demand will be high in a way that is not captured by our controls and firms invest accordingly, there will be positive correlation between cable quality and the unobserved elements of demand. In this case, we might overstate the role of cable quality such as capacity and number of cables. To address this concern, we exclude potentially endogenous cable variables from the set of instruments. In particular, $\mathbf{y}_{mt}^D$, $\mathbf{y}_{mt}^{D,\Delta}$, $\boldsymbol{y}_{rt}^Q$, and $\boldsymbol{y}_{rt}^{Q,\Delta}$ include year fixed effects, market or region fixed effects, and a time trend. That is, we take averages of $\xi_{mt}^D(\theta)$ across markets and time and also interact with time to form moments, and similarly for the other error terms.

These instruments satisfy the condition of exogeneity, but we also add additional instruments that are better targeted at identifying the endogenous coefficients. We assume that at the time of investment, firms cannot predict $\Delta\xi_{mt}^D(\theta)$ for future values of t. That is, firms cannot predict future demand innovations at the time of investment. This assumption is similar to Lee (2013) in the context of video games, which assumes that the introduction of new video games is orthogonal to future innovations in demand. To implement this, we include the one-period lagged values of capacity and the number of cables (K_{mt-1}^*, n_{mt-1}^*) in $\mathbf{y}_{mt}^{D,\Delta}$. We explore alternative assumptions in Section 5.1 and show that results are consistent with these assumptions about exogeneity.

4.3 Cable Investment

In this section, we use our model of cable investment to estimate γ_{jt}. In order to emphasize the dependence on parameters, we write $r_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t)$ as $r_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt})$.

We rely on moment inequalities derived from revealed preference, as in Pakes et al. (2015). Let $\boldsymbol{A}_t^{j,m,1}$ be the matrix $\boldsymbol{A}_t$ where the j, mth element is set to 1. We similarly define $\boldsymbol{A}_t^{j,m,0}$ for the case where the element is equal to 0. The incremental change in the revenue in market z to firm j from investing in market m is:

$$\Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) = r_{jzt}(\boldsymbol{A}_t^{j,m,1}, \mathcal{S}_t; \gamma_{jt}) - r_{jzt}(\boldsymbol{A}_t^{j,m,0}, \mathcal{S}_t; \gamma_{jt}).$$

Also, we define $\Delta\varepsilon_{jmt} = \varepsilon_{1jmt} - \varepsilon_{0jmt}$ and $\Delta\nu_{jmt} = \nu_{1jmt} - \nu_{0jmt}$. We define $\boldsymbol{a}_{mt}^{j,1}$ to be $\boldsymbol{a}_{mt}$ with element $a_{jmt} = 1$. Revealed preference tells us that if firm j invests in market m, t, it must be that:

$$\sum_{z \in \mathcal{M}} \Delta r_{jmzt}\left(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}\right) + \Delta\varepsilon_{jmt} - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + \Delta\nu_{jmt} \geq 0. \tag{11}$$

We take the expectation of Equation (11) conditional the firm's information set $\mathcal{J}_{jt}$ at observations for which we observe investment to obtain:

$$E\left[\sum_{z \in \mathcal{M}} \Delta r_{jmzt}\left(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}\right) - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + \Delta\nu_{jmt} \middle| a_{jmt} = 1, \mathcal{J}_{jt}\right] \geq 0$$

where $E[\Delta\varepsilon_{jmt}|\mathcal{J}_{jt}, a_{jmt}] = 0$ follows from the rational expectations assumption. Suppose that $\mathcal{Z}_{jt} \subsetneq \mathcal{J}_{jt}$ 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 $\boldsymbol{A}_t$. Then, we can apply the law of iterated expectations to derive:

$$E\left[\sum_{z \in \mathcal{M}} \Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + \Delta\nu_{jmt} \middle| a_{jmt} = 1, \mathcal{Z}_{jt}\right] \geq 0. \tag{12}$$

Similarly, we can define a separate moment for the value of investing for cases of no investment.

$$E\left[\sum_{z \in \mathcal{M}} \Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + \Delta\nu_{jmt} \middle| a_{jmt} = 0, \mathcal{Z}_{jt}\right] \leq 0. \tag{13}$$

Equations (12) and (13) still cannot be applied to the data as they depend on unobserved structural errors ($\Delta\nu_{jmt}$), which causes the well-known selection issue. That is, conditional on choices, the differences in the error terms are not mean zero. We follow the approaches in Pakes (2010) and Pakes et al. (2015) to address this issue and obtain two sets of inequalities.

The first set relies on the differencing approach following Pakes et al. (2015). We assume that $\Delta\nu_{jmt} = \Delta\nu_{j'mt}$ if $\mathcal{G}(j) = \mathcal{G}(j')$ where $\mathcal{G}(.)$ partitions firms into two groups: content providers and non-content providers. That is, all firms of the same type in a given market and year draw the same values for the difference of structural errors. We consider pairs of observations jmt and $j'mt$ for which $a_{jmt} = 1, a_{j'mt} = 0$ and $\mathcal{G}(j) = \mathcal{G}(j')$. Then, summing Equations (12) and (13) leads to:

$$E\Bigg[\sum_{z\in\mathcal{M}} \Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - \sum_{z\in\mathcal{M}} \Delta r_{j'mzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + c_{j'mt}\left(\boldsymbol{a}_{mt}^{j',1}\right) \Bigg| a_{jmt} = 1, a_{j'mt} = 0, \mathcal{Z}_{jt}, \mathcal{Z}_{j't}\Bigg] \geq 0. \tag{14}$$

We derive the following inequalities by interacting inequalities in (14) with positive-valued functions of $\mathcal{Z}_{jt}$ and $\mathcal{Z}_{j't}$, denoted as $h(\mathcal{Z}_{jt}, \mathcal{Z}_{j't})$:

$$\frac{1}{N^{a_{jmt}=1, a_{j'mt}=0}} \sum_{\{j,j',m,t|a_{jmt}=1, a_{j'mt}=0\}} \Bigg[\sum_{z\in\mathcal{M}} \Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - \sum_{z\in\mathcal{M}} \Delta r_{j'mzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + c_{j'mt}\left(\boldsymbol{a}_{mt}^{j',1}\right)\Bigg] h(\mathcal{Z}_{jt}, \mathcal{Z}_{j't}) \geq 0 \tag{15}$$

where $N^{a_{jmt}=1, a_{j'mt}=0}$ the number of pairs of firms j and j' in mt we use in the estimation, summed over all mt.

For the second set, we work with unconditional inequalities following Pakes (2010). The idea is that, for market-year pairs in which no firm decides to invest, we can form inequalities that do not condition on firms' actions by summing over $\Delta\nu_{jmt}$ across all firms. We let $\mathcal{H}$ denote the set of jmt that belong to mt in which no firm invests and let $N^{jmt\in\mathcal{H}}$ denote the number of firm-market-period observations belonging to $\mathcal{H}$. Then, we can sum inequalities in Equation (13) while assuming $E[\Delta\nu_{jmt}|Z_{jt}] = 0$ to obtain:

$$\frac{1}{N^{jmt\in\mathcal{H}}} \sum_{\{jmt\in\mathcal{H}\}} \Bigg[-\sum_{z\in\mathcal{M}} \Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) + c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right)\Bigg] h(\mathcal{Z}_{jt}) \geq 0, \tag{16}$$

given that the law of large numbers holds such that $\frac{1}{N^{jmt\in\mathcal{H}}}\sum_{jmt\in\mathcal{H}} \Delta\nu_{jmt} \xrightarrow{p} 0$.

Note that the inequality in Equation (16) does not condition on individual firms' investment decisions, and thus does not suffer from the selection of firms. One may still worry about the selection of markets, however, in that markets in which investment is unprofitable for unobserved reasons may be overrepresented. A very large portion of our observations (77%) fall into the market-period group with no investment, however, mitigating this concern. In order to further investigate this issue, we conduct a robustness check where we weight the observations in a way that all markets are equally represented in the inequalities construction following Pakes (2010) and Wollmann (2018) in Section 5.4.

We treat each cable as a separate investment opportunity for each firm and construct inequalities accordingly. That is, if we observe multiple cables being built in a single market-period, we allow firms to separately decide for each cable.

A final step to apply the inequalities is specifying $\tilde{K}(\boldsymbol{a}_{mt})$, the function that maps investment choices into capacity. That is, what happens if one of the firms changes its investment choice? Because we consider only single deviations from observed choices in our estimation approach, we need to specify how capacity would change if one firm changed its choice. Consider a counterfactual vector of choices $\boldsymbol{a}_{mt}$ which equals the observed vector at mt except for element j which has been switched to the alternative investment choice. We write the outcome from this case as:

$$\tilde{K}_j(\boldsymbol{a}_{mt}) = \begin{cases} 0 & \text{if } a_{zmt} = 0 \; \forall \; z \\ \hat{\tilde{K}}_{mt} & \text{if } a_{jmt} = 1 \text{ and } a_{zmt} = 0 \; \forall z \neq j \\ \tilde{K}^*_{mt} & \text{if } a_{zmt} = 1 \text{ for any } z \neq j. \end{cases} \tag{17}$$

The first line says that if no firms invest in the counterfactual set of choices, then no capacity is added. The second line covers the case when firm j is the only investor in the counterfactual set of choices, which results in a new cable that was not observed in the data. In this case, we impute added capacity from a regression, denoting the predicted value as $\hat{\tilde{K}}_{mt}$. The regression is described below in Section 5.2. The third line says that if any other firms invest, changing firm j's choice does not change the level of added capacity and we get the level of added capacity observed in the data, denoted as $\tilde{K}^*_{mt}$.

We focus on recovering the revenue function, particularly γ_{jt}, while calibrating the cost of investment using the construction cost data for each cable. An example of a paper using data on observed transaction costs to study investment in a capital-intensive industry is Jeon (2022). The cost calculation is described in Section 5.2.

We briefly discuss the intuition for the identification of the supply-side model. Recall that $q_{jmt}(\boldsymbol{A}_t, \mathcal{S}_t)$ denotes quantity served by firm j in mt. Denoting the incremental change in quantity served in market z induced by investing in market mt as $\Delta q_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t)$, then, we can rewrite Equation (15) as:

$$\frac{1}{N^{a_{jmt}=1, a_{j'mt}=0}} \sum_{\{j,j',m,t|a_{jmt}=1, a_{j'mt}=0\}} \Bigg[\gamma_{jt} \left(\sum_{z \in \mathcal{M}} \Delta q_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t) - \sum_{z \in \mathcal{M}} \Delta q_{j'mzt}(\boldsymbol{A}_t, \mathcal{S}_t) \right) \\ - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + c_{j'mt}\left(\boldsymbol{a}_{mt}^{j',1}\right) \Bigg] h(\mathcal{Z}_{jt}, \mathcal{Z}_{j't}) \geq 0. \tag{18}$$

In the case where firm j invests and j' does not, the term $\sum_{z \in \mathcal{M}} \Delta q_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t) - \sum_{z \in \mathcal{M}} \Delta q_{j'mzt}(\boldsymbol{A}_t, \mathcal{S}_t)$ is on average positive, which means that the inequality will yield a lower bound for γ_{jt}. Therefore, we are relying on the variation in the quantity gain from those that chose to invest *relative to* the quantity gain that chose not to invest in the same market and time period for the identification of the lower bound.

This intuition guides the choice of $j'mt$ in forming inequalities. Our goal is to make the comparison with firms that have the most similar set of cable capacities to firm j across markets to obtain the most informative lower bound. To do so, we choose $j'mt$ that minimizes the distance between $\sum_{z \in \mathcal{M}} \Delta q_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t)$ and $\sum_{z \in \mathcal{M}} \Delta q_{j'mzt}(\boldsymbol{A}_t, \mathcal{S}_t)$. In our baseline estimation in Section 5.3, we choose the five $j'mt$ that minimize the distance for each jmt. To test robustness, we vary the number of "nearest neighbors" from one to ten in Section 5.4 and find similar estimates.

We can similarly rewrite Equation (16), replacing $\sum_{z \in \mathcal{M}} \Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt})$ with $\gamma_{jt} \sum_{z \in \mathcal{M}} \Delta q_{jmzt}(\boldsymbol{A}_t, S_t)$. Since the term $\sum_{z \in \mathcal{M}} \Delta q_{jmzt}(\boldsymbol{A}_t, S_t)$ is on average positive, this inequality will yield an upper bound for γ_{jt}. And an increase in the magnitude of the quantity term, $\sum_{z \in \mathcal{M}} \Delta q_{jmzt}(\boldsymbol{A}_t, S_t)$, or a decrease in the cost term, $c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right)$, will shift the upper bound downward. Intuitively, if firms choose not to invest even when there are large quantity gains from doing so (or even when the cost of investment is low), we can infer that the profit from investing must be low.

Our framework is static although cable investments are inherently dynamic decisions made by forward-looking firms. Dynamics are difficult to handle computationally because of the large state space in our model. Due to the networked nature of the cable market, a firm's decision depends not only on the state of the market in which the firm is considering investment, but also on the states of all other markets. Moreover, modeling the fast-changing landscape of the industry as a stationary setting is a challenge. To provide robustness on dynamics, we consider in Appendix F a dynamic model under the assumptions of perfect foresight and commitment. A factor that lends support to these assumptions in our market is that firms require government approval to land their cable so they typically know their own and rival cable construction years in advance. These assumptions are computationally attractive because we can still derive moment inequalities from a deviation by a single firm without having to solve for how rivals adjust their behavior, similar to how we derive moment inequalities above.

5 Results

5.1 Results for Demand Estimation

Table 2 presents the demand-side estimates for alternative choices of moments and instruments. Our main results are in Column 1. Panel A shows the demand coefficients, θ^d. The estimates suggest that a one-percent increase in the sum of the log of GDP of a pair of countries involved in trade would result in an approximately 1.26 percent increase in used bandwidth, and a one-percent increase in the sum of the log of population would result in a -0.30 percent decrease in used bandwidth. Thus, an increase in both GDP and population (an increase in GDP per capita) would lead to an increase in used bandwidth.

As described in Section 2.3, we include dummies for our three main time periods: before 2010, 2010-2017, and 2018-2020. We find small negative coefficients for the two later periods, but also a strong time trend coefficient, so overall, demand is growing over time. Note that we do not include distance between countries as an explanatory variable, which would be standard in a gravity regression. Instead, we include market fixed effects, which capture distance and we found valuable in matching the data.

Panel B presents the path coefficients, θ^δ. Consistent with our expectations, path length has a negative effect, while path capacity and the number of cables on a path 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

Table 2: Demand model estimates with alternative choices of moments and instruments

	(1)	(2)	(3)	(4)
FD Moments Included	Y	N	Y	Y
Lagged cable vars IVs for FD	Y	NA	N	Y
Lagged cable vars IVs for Levels	N	N	N	Y
Panel A. Demand for data: (θ^d)				
Constant	-182.32	-226.65	-155.18	-218.59
	(71.33)	(61.29)	(261.90)	(11.33)
GDP	1.26	3.92	0.70	1.40
	(0.12)	(1.39)	(0.05)	(0.03)
Population	-0.30	-1.85	-0.92	0.97
	(0.11)	(0.91)	(0.03)	(0.06)
Time trend	0.36	0.49	0.36	0.23
	(0.01)	(0.18)	(0.01)	(0.01)
2010-2017	-0.47	-1.00	-0.46	-0.58
	(0.03)	(0.35)	(0.03)	(0.02)
2018-2020	-0.80	-0.02	-0.70	-0.67
	(0.06)	(0.42)	(0.06)	(0.03)
Market fixed effects	Y	Y	Y	Y
Panel B. Cable usage: (θ^δ)				
Constant	126.15	93.24	142.01	89.56
	(70.21)	(93.96)	(262.64)	(12.56)
Path length	-1.82	-2.23	-2.56	-0.02
	(0.09)	(1.21)	(0.16)	(0.04)
Potential capacity	1.16	0.93	1.07	3.13
	(0.04)	(0.13)	(0.06)	(0.05)
Number of cables	0.12	0.21	0.07	0.01
	(0.01)	(0.10)	(0.01)	(0.01)

Notes: Panel A reports estimates for θ^d, the parameters that govern how country-pair characteristics (x_{ckt}) affect demand for data transmissions between the two countries. Panel B reports estimates for θ^δ, the parameters that govern how cable features of a path 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 vector x_{ckt} also includes a linear time trend, indicators for the 2010-2017, and 2018-2020 periods, and market fixed effects. The 2002-2009 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. "FD Moments included" refers to whether the first-difference moments ($\boldsymbol{m}^{D,\Delta}(\theta)$ and $\boldsymbol{m}^{Q,\Delta}(\theta)$) are included. "Lagged cable vars IV for FD" refers to whether lagged values of capacity and number of cables are included in the instrument vector for the first difference moments (i.e, included in $\boldsymbol{Y}^{D,\Delta}$ and $\boldsymbol{Y}^{Q,\Delta}$). "Lagged cable vars IVs for Levels" refers to whether these variables are included in the instrument vector for the moments in levels ($\boldsymbol{Y}^D$ and $\boldsymbol{Y}^Q$). The sample includes 179,778 country-pair-year observations and 1,710,964 country-pair-path-year observations. All specifications from Columns 1 to 4 use 399 market-year and 133 island-year observations of used bandwidth. Column 1 uses 188 moments in total, column 2 uses 62, Column 3 uses 122, and Column 4 uses 254.

sometimes value having more options on a path in case one cable malfunctions, as in Caoui and Steck (2025). The coefficients in Panel B determine how much usage goes over the shortest path (the direct path between two regions if a direct path exists) versus more circuitous paths. Thus, these coefficients govern how much an investment in one market affects data flows in other markets, more fully explored below.

We also consider a nested logit model in which all paths are in one nest and the outside option has its own nest. However, we found (in unreported results) that the inclusive parameter was very close to 1, i.e., the standard logit model fit the data well. In our related paper (Jeon and Rysman, 2024), we explore a variety of additional variables such as whether the two countries share a common language or have a trade agreement.

Column 1 includes the moments both in levels and first differences and applies lagged capacity and number of cables to the first-difference moments. We consider alternative setups in Columns 2 to 4. Column 2 drops the first difference moments. For the coefficient on capacity, Column 2 finds a somewhat smaller coefficient than Column 1 with a much larger standard error. In fact, there is a larger standard error for all of the θ^δ parameters. This points to the importance of the first difference moments in obtaining power over these parameters.

Column 3 includes the first difference moments but does not apply lagged capacity and number of cables as instruments. The coefficient on capacity is somewhat smaller than in Column 1 with a standard error that is about 65% larger. That is consistent with our view that the instruments in Column 3 are exogenous but lack a variable with strong predictive power over capacity. In Column 4, we apply lagged capacity and the lagged number of cables as instruments to both sets of moments. We believe that applying these instruments to the moment in levels raises an endogeneity problem since lagged capacity may respond to the expected level of demand. Indeed, we see that the coefficient on capacity is much larger in Column 4, as we would expect if the instrument does not resolve the endogeneity problem. Overall, the results in Table 2 show coefficient changes that we would expect under our interpretation of the endogeneity problem.

We highlight the network nature of the market by examining which end-point regions drive the data flows in a given market. For example, Panel A of Table A.1 shows that 86% of the usage in cables in the Europe-North America market in 2020 can be attributed to data flows 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. An alternative approach is to pick a set of endpoints and consider how much usage is allocated to different paths. Panel B presents this calculation for Europe-

South America, where we see that 91% goes through the Europe-North America-South America path. It is still a very small share in the Europe-North America market because the Europe-South America endpoints account for less than 1% of the total data flows between Europe and North America, which gives a sense of the vastly different scales across different markets.

5.2 Capacity and Cost Imputation

As described in Section 4.3, we impute capacity in markets where no investment is observed when we calculate counterfactual investment choices. For this calculation, we specify a simple linear model:

$$\ln(\tilde{K}_{mt}) = \beta_0 + \beta_1 t + \sum_{r=1}^{R} \alpha_r \mathbb{1}\{r \texttt{ in market } m\} + \omega_{mt}.$$

In this model, new capacity is a function of a time trend and region fixed effects α_r, where α_r applies for both regions in a given market. That is, if regions 1 and 2 are in market 3, then α_1 and α_2 apply to the determination of $\tilde{K}_{3t}$. We estimate this model via OLS.

We also need cost for each investment choice. We observe cost for cables in our data for most but not all observed cables. We are not aware of a selection issue for which cables we observe cost. We specify a model similar to the one for new capacity: a linear model of the log of cost on a time trend, cable capacity, and cable length, estimated by OLS. We do not include region-fixed effects in the cost regression because they add very little; R^2 is quite high already.

The full results appear in Table A.2. To summarize the results, for the capacity regression, we find a coefficient on the time trend of 0.144, which implies an annual growth rate over 15%, consistent with the massive growth in cable capacity and internet usage we saw in descriptive statistics.[16] Note that when including log distance in this regression, it has an insignificant coefficient and the other parameters change little. Presumably, distance is well-captured by the region fixed effects.

For the cost regression, we find a coefficient on the time trend of -0.028. That is, the cost of installation falls almost 3% per year holding all else equal. The coefficients on the logged length and capacity variables imply that doubling the length of a cable

[16] Because the dependent variable is logged and time is in levels, the percentage increase from an additional unit of time is $\exp(\beta_1) - 1$. Our result is $\exp(0.144) - 1 = 15.49\%$.

almost doubles the cost, whereas doubling capacity increases the cost by a little under 14%. That is consistent with laying the cable (or perhaps manufacturing the cable) being a major determinant of cost whereas increasing the thickness of a cable has a more moderate effect on cost. Over time, cost falls for a given capacity, but the capacity of cables increases. The overall effect is such that costs for a cable fall slowly over time even accounting for the increase in capacity, but the effect is small.[17] Note that we use costs in nominal dollars rather than real dollars due to the lack of a price index targeted for this industry. Naturally, deflating by a standard price index would imply more substantial cost declines over time.

When we need to input the value of cost or capacity, we use the expectation of cost or capacity calculated from these regressions. We use the method of Duan (1983) to convert from the log regression to the expectation of the level of each variable.

5.3 Results for Investment Model Estimation

Table 3: Supply model estimates

	Carriers	Content providers
2002-2009	[2674.3 , 3517.9]	NA
2010-2017	[22.3 , 98.8]	[11.3 , 96.0]
2018-2020	[12.0 , 15.6]	[5.0 , 18.8]

Notes: We report the 95% confidence set for the markup parameter, γ_{jt}, capturing the profit (measured in millions of US dollars) from an increase in the quantity served by the firm by one Tbps. The estimates are missing for the 2002-2009 period for content providers due to the lack of observed investment episodes. The sample includes 56,935 firm-cable-market-year-level observations for carriers and 1,180 for content providers.

We allow the markup parameter γ_{jt} to vary by firm type (content providers vs. non-content providers) and across three time periods. Our inference method is based on Cox and Shi (2023) with detailed steps provided in Appendix D. Table 3 presents the estimates of γ_{jt}. In the 2002-2009 period, we recover a 95% confidence interval of [2674.3 , 3517.9] for carriers, which implies that an increase in a firm's quantity by 1 Tbps would lead to a profit increase ranging from 2.7 to 3.5 billion dollars. By 2018-2020, this number decreases to a range of 12 to 15.6 million dollars. This steep decline in the markup is consistent with alternative data on bandwidth prices.[18] The

[17]The indirect effect of time through capacity can be computed as 0.144 (time trend coefficient in capacity) $\times$ 0.137 (capacity coefficient in cost) = 0.02. Since the direct effect of time on cost is -0.028, the full effect of time is 0.02-0.028 =-0.008.

[18]TeleGeography provides data on median monthly lease prices for 10 Gbps wavelength for a

confidence interval for content providers mostly overlaps with that of carriers. It ranges from 11.3 to 96 million dollars in 2010-2017, and 5 to 18.8 million dollars.

5.4 Robustness

We consider a number of alternative specifications for the markup parameter. Results appear in Table A.4. We first estimate the parameter under two alternative assumptions on the structural error term, ν_{ajmt}. The first approach is to assume $\Delta\nu_{jmt} = 0$ for all jmt, that is, there is no selection on unobservables. Alternatively, we can make a weaker conditional independence assumption: $E\left[\Delta\nu_{jmt}|a_{jmt}, \mathcal{Z}_{jt}\right] = 0$; similar assumptions are found in Holmes (2011) and Ho (2009), for example. Under either of these assumptions, the term $\Delta\nu_{jmt}$ drops out of the inequalities in (12) and (13). The estimates under this no selection approach are shown in Panel A of Table A.4. We find that the confidence intervals under this approach are narrower than our baseline confidence intervals, which is not surprising given that we are making a stronger assumption on the unobserved term. But the results in Panel A are very close to our main results.

In Panel B, we examine whether the lower bounds are sensitive to the choice of $j'mt$ in forming inequalities in Equation (18). In the baseline, we form five pairs for each jmt, picking the five that minimize the distance between $\Delta q_{jmzt}(\boldsymbol{A}_t, S_t)$ and $\Delta q_{j'mzt}(\boldsymbol{A}_t, S_t)$. We vary this number from one to ten and find that the bounds are robust to this change.

In Panel C, we conduct another robustness check with the inequalities construction. One worry in using market-periods in which no investment is observed to form inequalities in equation (16) is that certain markets that have unprofitable investment opportunities for unobserved reasons may be over-represented. Following the approach in Pakes (2010) and Wollmann (2018), we adjust weights to observations such that each market is represented equally.[19] The bounds are slightly wider, but

limited set of routes from 2015 to 2019. We construct the average across all available interregional routes and find that it fell by 29% yearly on average in this period. As a comparison, we compute the compound annual growth rate of our estimated markup term using the midpoint in each of the three time periods and the midpoints in the confidence intervals for the carriers. For example, from 2002-2009 to 2010-2017, we assume that γ_{jt} fell from \$3.1 billion per Tbps to \$60.6 million per Tbps over 8 years. The rate is -39% from 2002-2009 to 2010-2017 and -24% from 2010-2017 to 2018-2020. Additionally, we show in Table A.3 that internet transit prices from 1998 to 2015 show an even faster rate of decline according to Norton (2014).

[19]We implement this procedure in the following way. Define $\mathcal{H}_m$ as the set of observations from market j belonging to periods with no investment. Then, we multiply each observation with $\frac{N^{\text{total}}}{N^{jmt\in\mathcal{H}_m}}$

robust to this change.

Another concern is that many of the owners in our sample appear only once and we would not want our estimates to be entirely driven by this group. Moreover, we would like to capture whether there is important heterogeneity in the profit margin across firms due to market power, cost differences, or other reasons. To address these issues, we separate owners into two groups: One group for the 27 major carriers (including all content providers) that meet a set of criteria involving the total owned capacity, the number of active markets, etc., and one group for the rest of the 170 carriers that we call 'small carriers.' We separately estimate bounds on γ_{jt} for these two groups. We rely on the non-differencing approach that assumes no selection on unobservables because our baseline differencing approach leads to an empty confidence interval for certain time periods. As shown in Panel D, the estimates are similar across these two groups and overlap significantly in all three periods.

In Panel E, we allow the capacity of a new cable to depend on the number of investors. In particular, we assume that adding or subtracting an investor raises or lowers capacity by 5% and rely on the non-differencing approach. We find that the bounds are slightly wider but remain similar.

We now turn to estimates of our dynamic model, as described in Appendix F. Relative to our static model, we expect to find values of γ_{jt} that are much lower. That is because in the static model, the estimates of γ_{jt} represent firms' perceived discounted sum of flow profits per unit of quantity served, assuming that the firms believe that the quantity stays constant in the future. More broadly, allowing firms to believe that the quantity of data flows will change in the future, we can interpret the estimates to incorporate agents' beliefs about changes in future quantities in addition to future profit margins. Thus, we are more interested in the relative values we estimate than the levels. As shown in Table F.10, the estimates show similar patterns to the baseline estimates. Taking the midpoints in each of the three time periods and the midpoints in the confidence intervals for the carriers, we compute the compound annual growth rate to be -36% from 2002-2009 to 2010-2017, and -15% from 2010-2017 to 2018-2020 for carriers. These rates are slightly lower compared to those from the static model, which can be due to the fact that the dynamic model takes into

when computing moments in Equation (16) where $N^{\text{total}} = J \times T$ is the total number of observations (jmt) in each market and $N^{jmt \in \mathcal{H}_m}$ is the number of observations (jmt) in market m belonging to periods in which there is no investment.

account the future growth in demand.

Lastly, a potential concern with our approach is that we ignore geopolitical issues that might lead to limitations in investment options, such as potentially limiting US firms from landing cables in China, or vice versa. It would be straightforward to incorporate this into our model, for example, by eliminating the choices of Chinese firms to not invest in North American markets from our construction of moment inequalities. However, in our data, we see numerous examples of US firms investing in the East Asia markets and Chinese firms investing in the North American markets, as well as co-investment by Chinese and US firms. The same is true if we focus on Europe and European firms. These government restrictions are difficult to evaluate because they are not typically implemented through explicit policies, but rather through ad hoc decisions by bodies like Team Telecom.[20] In any case, we do not see ownership and investment patterns consistent with restrictive policies on ownership decision-making and we do not further explore this issue here.

6 Counterfactuals

6.1 Understanding the Growth in Data Flows

Our model allows us to study which factors have contributed to the growth in total data flows. We decompose the observed growth into the part driven by (a) changes in demand for data flows, 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. This exercise uses only the estimates from demand estimation (Table 2), not the supply side model.

We first compute the total data flows for the world based on our model from 2002 to 2020 as: $\sum_{c=1,\dots,C-1}\sum_{k=c+1,\dots,C}\hat{d}_{ckt}$ where $\hat{d}_{ckt}$ is the predicted data flows for country pair ck based on Equation (7). 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} (for simplicity, we now refer to $v_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ as v_{ckt}) to

[20] It may be that the restrictions are more focused on equipment manufacturers (e.g., a cable using telecom equipment produced by Huawei) than on ownership. Industry statements support our result. For instance, Lipman and Tanner (2010) writes about cable construction: "Foreign persons, however, should not be dissuaded from entering the U.S. market. Nearly all projects can ultimately receive approval through proper structuring to address basic security concerns." (pg. 1)

Figure 4: The decomposition of the growth in data flows

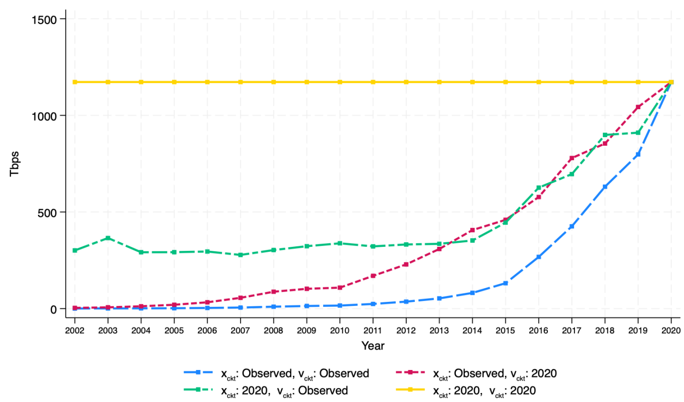

Notes: This figure plots the total data flows computed based on our model estimates, holding country-pair characteristics (x_{ckt}) and quality of connection (v_{ckt}) at the observed levels. This figure also plots data flows under the following three counterfactuals: (i) holding x_{ckt} and v_{ckt}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

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 4. 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 increase total data flows by approximately 10 times. Endowing the world with the 2020 level of x_{ckt} while fixing the cable network at the 2002 level, however, would have a much larger effect, bringing the data flows to over one-quarter of the 2020 level. The contribution of the quality of cable connections grows through the sample period. For example, in 2013, improving the cable connections to the 2020 level while fixing x_{ckt} at the observed level would increase the data flows by 6 times, or bring it up to 26% of the 2020 level. This is comparable to endowing the world with the 2020 level of x_{ckt}, which also increases the data flows by 6 times. From 2013 to 2018, the demographic characteristics (and the time trend) and the quality of cable connections continue to have comparable contributions to the growth in data flows. In 2019, the role of the cable network surpasses that of the other drivers. Overall, we find that the cable network was an important contributor to the growth in internet usage.

6.2 Inefficiencies in the Subsea Cable Network

The network nature of the market along with the high level of geographical concentration in the cable network suggests the possibility that cables are misallocated to markets that generate relatively low social values from these cables. We characterize the inefficiencies in the cable network and examine their sources. We first solve the planner's problem. We take the cable network at the end of 2020 and consider 15 markets that have positive capacity (i.e., at least one cable) in our sample period. We consider all combinations of these markets to build a cable in and identify the combination that would yield the highest increase in total welfare. We calibrate the capacity of these counterfactual cables based on the method described in Section 4.3 and rely on the method described in Appendix Section E for computing consumer surplus.[21]

The map of the efficient allocation appears in Figure 5, with the red lines indicating markets the planner would find optimal to invest in and the blue lines indicating the rest of the markets. The thickness of the lines reflects the total capacities in the market. We find that the planner would build a new cable in 8 out of the 15 markets. This includes not only markets that have high demand and existing capacity such as Europe-North America and North America-South America, but also markets that have received limited investment such as East Asia-Sub-Saharan Africa and South Asia-Sub-Saharan Africa.

We examine whether private incentives are misaligned with social payoffs, which would lead to spatial misallocation. To do so, we consider a de novo entrant who does not own any existing cable constructing a new cable in one market at the end of 2020. We compare the profits the entrant would earn and the change in total surplus that the new cable would generate. Figure 6 plots the results for each market.[22] We

[21]Computing consumer surplus is not straightforward due to the lack of detailed information on the pricing of cable capacity leasing. We rely on back-of-the-envelope calculations based on limited pricing data from TeleGeography and a set of assumptions about the demand function, particularly constant elasticity of demand making use of γ_{jt} as the measure of the mark-up. See Appendix Section E for more details.

[22]In evaluating our results on entrant profits, it is helpful to recognize how our model differs from a free entry model, which would imply that profits to an additional entrant are negative in every market. Our model instead implies that in markets with no investment, entrants must have negative profits on average across markets, which means that it is possible to predict positive entrant profits even in markets in which no investment is observed. Moreover, entry in our model leads to an exogenous increase in capacity, and that may not exhaust profitable opportunities in the market. Our results are broadly consistent with this. Early years have relatively little investment and we

Figure 5: The efficient cable network

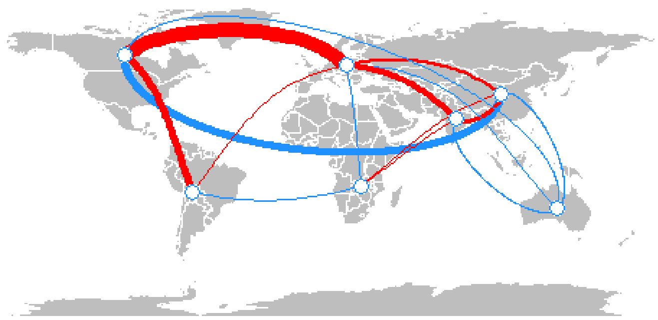

Notes: This figure shows the map of the subsea cable network in which the thickness of the lines reflects the total cable capacity of the market. The red line indicates that the planner would find it optimal to add a cable in the market at the end of 2020. The blue lines mark markets with existing cables in which the planner would choose not to add a new cable. The results are the same regardless of whether we use the lower or upper bound of our estimates.

use results based on lower bounds in the rest of the discussion, but the results based on upper bounds are close due to the tightness of the bounds in this period.

In general, there is a positive correlation between the entrant profit and the change in total surplus across markets. For example, in 2020, building in high-demand markets like Europe-North America would lead to a large increase in both the entrant profit and total surplus, suggesting that private incentives are aligned well with social values in these markets. Similarly, in many of the markets in which welfare gain is small or negative, an entrant would also have no incentive to enter. However, there are several markets in which entrants are not incentivized to invest even though doing so would generate significant welfare gains. This mismatch happens in East Asia-Sub-Saharan Africa, Europe-Sub-Saharan Africa, and Europe-South Asia. Furthermore, among the markets in which the cable entry would generate a large welfare gain, the entrant profit (i.e., private value) and the total surplus gain (i.e., social value) are not necessarily ranked in the same way. For example, even though the welfare gain is higher in East Asia-South Asia and Europe-South America, the entrant profit is substantially higher in East Asia-North America.

There are two main externalities in our model that can lead to misaligned incentives: business stealing and spillovers across markets. It is evident from Figure 6

find entrant profits are negative in most markets. But in later years, our model finds that markets with repeated investment, such as North America-Europe, start to show positive entrant profit. In 2020, entrant profit is positive on average across markets.

Figure 6: Profits and welfare changes from new cable entry

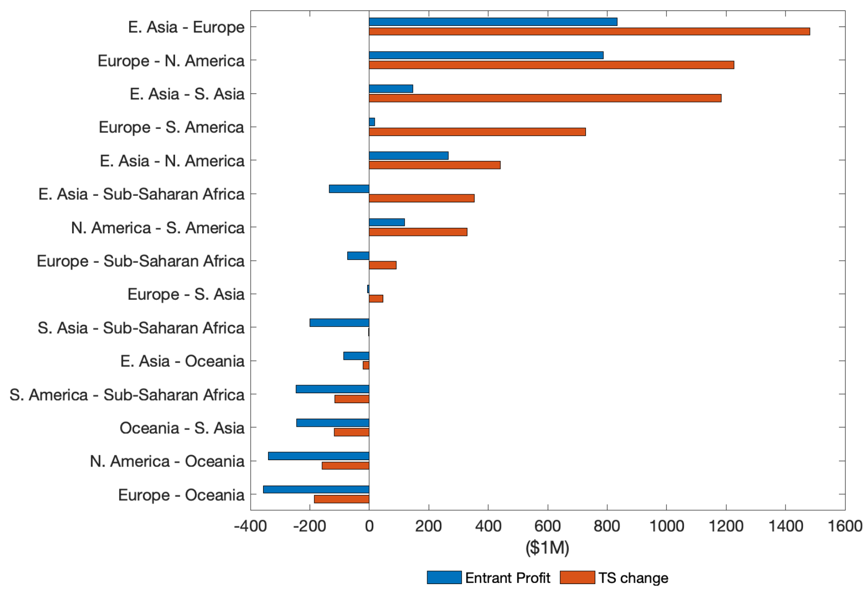

Notes: We consider an entry of a new cable by a de novo entrant in each market separately. For each market, the first bar represents the entrant profit and the second bar the change in total surplus. The markets are ordered by the total surplus change. All figures are based on the lower bounds of the estimates and are in millions of US dollars.

that business stealing is not driving the result. If business stealing were important, we would observe markets where entrant profit was positive and social surplus was negative. However, there are no such cases, or even cases in which entrant profit is greater than social surplus. Instead, misalignment is due to the networked nature of the market.

To further understand our results, we decompose the effect of cable entry into (i) entrant profit; (ii) change in incumbent profits in the investment market; (iii) change in incumbent profits in the rest of the world; (iv) change in consumer surplus in the investment market (that is, in countries in the regions connected by the market); and (v) change in consumer surplus in the rest of the world. Note that the direction of the effects on incumbent profits is theoretically ambiguous. A new cable transfers market shares from incumbents to the entrant in the investment market. An increase in the capacity in that market also diverts data flows away from the rest of the world, lowering profits there as well. At the same time, the increased capacity increases demand for data flows, resulting in market expansion, typically in multiple markets.

Tables A.5 to A.8 present full results for all markets across different years. Fo-

Figure 7: The decomposition of the effects of cable entry

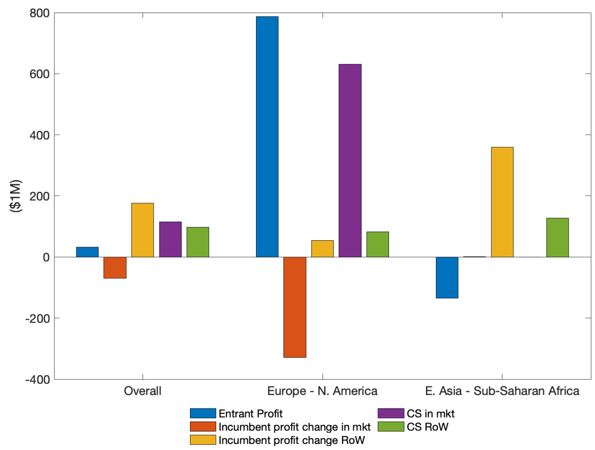

Notes: We consider an entry of a new cable by a de novo entrant at the end of 2020 in three selected markets: Europe - North America, Europe - South America, and East Asia-Sub-Saharan Africa. For each market, the first bar represents the entrant profit; the second bar, the change in incumbent profits in the market in which the investment happens; the third bar, the change in the rest of the world; the fourth bar, the change in consumer surplus in the investment market; and the last bar, the change in consumer surplus in the rest of the world. All figures are based on the lower bounds of the estimates.

cusing on the results for 2020 in Figure 7, we find that, on average, incumbents collectively lose 69 million dollars in profit in the market in which a new cable arrives. This suggests that the business stealing effect dominates the market expansion effect in the investment market. In contrast, outside the investment market, incumbents benefit from market expansion, collectively earning a profit gain of 176 million dollars. The entrant earns a profit of 32 million dollars on average. The consumer surplus increases by 115 million dollars and 98 dollars in the investment market and the rest of the world, respectively.

We now turn to the results for two specific markets: Europe-North America (EUR-NAM) and East Asia-Sub-Saharan Africa (EAS-SSA).[23] In both these markets, a new cable substantially increases total welfare (by approximately 1.2 billion dollars in EUR-NAM and 0.4 billion dollars in EAS-SSA). Interestingly, these two markets have very different cable profiles as of 2020. EUR-NAM has the highest capacity (1,118 Tbps) and the highest number of cables (21) among all markets. In contrast,

[23]The first represents markets with large capacity such as East Asia-Europe and East Asia-North America, and the latter small-capacity markets. As shown in Table A.8, a number of small markets such as Europe-South America and South Asia-Sub-Saharan Africa show similar results to East Asia-Sub-Saharan Africa, for example.

EAS-SSA has one viable cable, which has a capacity of 1.7 Tbps. This suggests that most communications between the two regions are carried through indirect paths, for example, through East Asia-Europe and Europe-Sub-Saharan Africa.

A new cable investment generates contrasting outcomes in these two markets. Incumbent profits decrease by 340 million dollars in EUR-NAM, while increasing by a small amount (1.5 million dollars) in EAS-SSA. Incumbents also experience disproportionately higher gains in the rest of the world from the investment in EAS-SSA relative to EUR-NAM (360 million dollars vs. 54 million dollars). This is because the investment in EAS-SSA results in greater spillovers to other markets relative to within-market market expansion. An entrant gains a high profit of 787 million dollars from EUR-NAM. In EAS-SSA, however, it would require a subsidy of 134 million dollars to incentivize an entrant, since it is unable to capture the network externalities. The consumer surplus gain also reflects this asymmetry. From the new cable in EUR-NAM, consumer surplus increases more within the investment market relative to the rest of the world (631 million vs. 82 million dollars). In contrast, for EAS-SSA, almost all benefits to consumers fall outside the investment market.

Figure 8: The revenue change from entry by firm

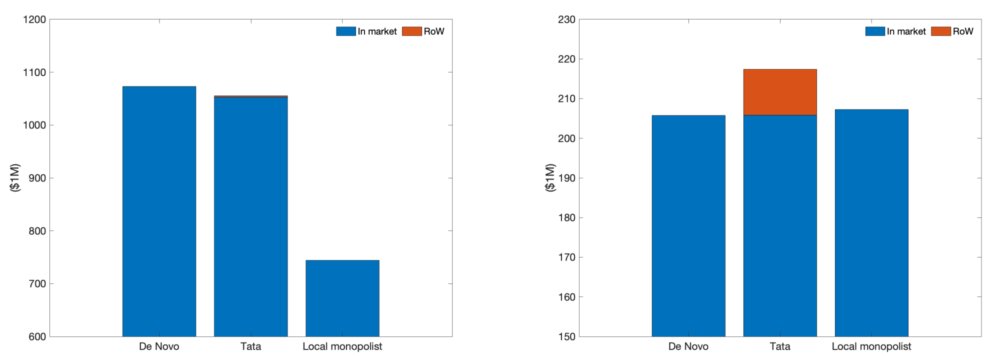

(a) EU-N. America (b) E. Asia-SSA

Notes: We consider an entry by three different firms (a de novo entrant, Tata Communications, and a local monopolist) at the end of 2020 into Europe - North America and East Asia-Sub-Saharan Africa. For each market, the first bar represents the change in the revenue for a de novo entrant, the second bar, for Tata Communications, and the last bar for a local monopolist. The revenue does not include the cost of investment. We break the change into the part that arises from the market in which the investment happens (in blue) and all other markets (in red). All figures are based on the lower bounds of the estimates.

Lastly, we investigate whether there is important firm heterogeneity. Which type of firm enjoys higher profit gains from a cable investment depends on the size of market expansion and the incumbent's ownership of cables across the entire network. For example, if an incumbent that owns some capacity in a market constructs a new cable in the market, it would receive an additional gain from market expansion on the existing capacity it already owns, but also incur a loss from cannibalization on the existing capacity.

We study this by comparing outcomes for three different firms: a de novo entrant, Tata Communications, and a local monopolist. The local monopolist owns all the cables in the market where the investment happens, but no cables elsewhere. Therefore, unlike the de novo entrant, the local monopolist would not be subject to the negative externality it imposes on incumbents (business stealing). Tata has a global presence in the cable network, owning 220 Tbps in 2020 across all 7 regions in our data, thus would internalize some of the network spillovers, unlike the de novo entrant or local monopolist.

Figure 8(a) plots results for the investment in EUR-NAM. The cannibalization effect dominates in the investment market, so Tata would earn less profit than a de novo entrant, but Tata benefits from out-of-market spillovers. Overall, Tata would profit less from this investment. In a market such as EAS-SSA where no incumbent owns much capacity, however, carriers like Tata that own large capacity outside the market gain a high profit. Therefore, although the difference is small, it would generally require a smaller subsidy to incentivize larger carriers like Tata than a de novo entrant to build a new cable in these markets.

The fact that the revenue is significantly higher for a de novo entrant than a local monopolist for EUR-NAM shows that there are substantial business stealing incentives in this market. Although we saw in Figure 6 that business stealing does not generate misalignment between entrant profits and total surplus, it is still the case that business stealing has a large magnitude in some cases.[24]

Overall, the results highlight that private incentives do not always align with what is socially optimal. Markets in which cable investment would yield high social gains do not necessarily guarantee high private gains, suggesting that there is a substantial misallocation of investment. An optimal network would require building new cables not only in the largest markets, but also in some of the most underserved markets

[24] Figure A.2 shows profits to a de novo entrant and a local monopolist from a new cable investment for all markets. We find that there is no market with a positive total surplus change such that a local monopolist would find the investment profitable, but not a de novo entrant.

such as Europe-South America and East Asia-Sub-Saharan Africa. However, because of the global nature of the externality, even individual national governments are not sufficiently incentivized to address the difference between social and private gain. A national government or 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. Our results point to the benefits of global coordination for investment support, although it may be difficult to organize effectively.

7 Conclusion

Undersea internet cables are a critical piece of communication infrastructure, carrying more than 99% of overseas 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 a model of investment and use of undersea internet cables. Potential investors choose simultaneously whether to invest in each market (a pair of regions). Investment leads to an increase in capacity in that market. If multiple firms choose to invest in a given market, they become collaborators. Country pairs exchange data over the resulting network, generating profits for the firms, where cable capacity and length affect outcomes. Our model emphasizes two critical phenomena. First, business stealing opportunities attract investment even if they offer little improvement in total surplus. Second, investment in one market affects traffic flows throughout the world because of the networked nature of Internet communication. Both because of the business stealing externality within a market and because of spillovers to other markets, private and public payoffs to investment may deviate substantially.

We structurally estimate the model using new data on the subsea cable industry. We utilize a moments approach to estimate the model of demand flows, based in part on a reduced-form approximation to the behavior of routers directing traffic over the network. Our model infers country-to-country demand even though we observe flows between regions, not the end points of those flows. We estimate the model of investment choice using moment inequalities with a variety of robustness checks. We find that country-level variables, such as GDP and GDP per capita are strong predictors of demand for used bandwidth. Also, cable path characteristics such as capacity and distance are important drivers of used bandwidth on individual cables. We find that content providers make similar revenue to carriers, and that revenue per unit falls substantially over time.

We use our results to make two points. First, construction of the cable network contributes significantly to the growth of internet usage over and above any secular trends or changing demographics. Second, we identify significant examples of mismatch between incentives to invest and the resulting total surplus. We find that in large markets with high existing cable capacity, an investment in a new cable tends to be socially valuable, and also highly profitable to entrants. Business stealing incentives are substantial, but mostly in markets where a new cable generates large social values, therefore are not critical in creating mismatches between private and social incentives. We find that network spillovers play a more important role in creating misaligned incentives. In several smaller markets with little existing capacity, when a new cable enters, a high share of market expansion falls outside the market, not providing enough incentives for entrants to invest, despite large social values. Even though incumbents like Tata Communications that own a large network of cables would internalize more of the network externalities, their ownership is not large enough to find it profitable to enter these markets.

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, Victor and Chun-Yu Ho, "A Dynamic Game of Airline Network Competition: Hub-and-spoke Networks and Entry Deterrence," *International Journal of Industrial Organization*, 2010, *28*, 377–382.

Allen, Treb and Costas Arkolakis, "The Welfare Effects of Transportation InfrastructureImprovements," *Review of Economic Studies*, 2022, *89*, 2911–2957.

Andrews, Donald W. K. and Panle Jia Barwick, "Inference for Parameters Defined by Moment Inequalities: A Recommended Moment Selection Procedure.," *Econometrica*, 2012, *80* (6), 2805–2826.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin, "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, September 2007, *75* (5), 1331–1370.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papgeorgiou, "Geography, Transportation, and Endogenous Trade Costs," *Econometrica*, 2020, *88*, 657–691.

Canay, Ivan A., Gaston Illanes, and Amilcar Velez, “A User’s Guide for Inference in Models Defined by Moment Inequalities,” June 2023.

Caoui, El Hadi and Andrew Steck, “Network Diversity, Market Entry, and the Global Internet Backbone,” *SSRN Electronic Journal*, 2025.

Cariolle, Joel, “International connectivity and the digital divide in Sub-Saharan Africa,” *Information Economics and Policy*, 2021, *55*, 100901.

Cariolle, Joël, Georges Vivien Houngbonon, Tarna Silue, and Davide Strusani, “The Impact of Submarine Cables on Internet Access Price, and the Role of Competition and Regulation,” Technical Report, The World Bank 2024.

Clark, David and Sara Wedeman, “Measurement of Internet access latency: A cross-dataset comparison,” Unpublished manuscript, CSAIL 2024.

Clark, David D., *Designing an Internet*, Cambridge, MA: MIT Press, 2018.

Cox, Gregory and Xiaoxia Shi, “Simple Adaptive Size-Exact Testing for Full-Vector and Subvector Inference in Moment Inequality Models,” *Review of Economic Studies*, 2023, *90*, 201–228.

CRS, “Undersea Telecommunication Cables: Technology Overview and Issues for Congress,” Technical Report September 2022.

Dickstein, Michael J and Eduardo Morales, “What Do Exporters Know?,” *The Quarterly Journal of Economics*, November 2018, *133* (4), 1753–1801.

Duan, N., “Smearing Estimate: A Nonparametric Retransformation Method,” *Journal of the American Statistical Association*, 1983, *78*, :605–610.

Eichengreen, Barry, Romain Lafarguette, Arnaud Mehl, and Massimo Ferrari Minesso, “Technology and the Geography of the Foreign Exchange Market,” *Journal of International Money and Finance*, 2023, *131*, 102802.

Eizenberg, A., “Upstream Innovation and Product Variety in the United States Home PC Market,” *The Review of Economic Studies*, 2014, *81*, 1003–1045.

Elliott, Jonathan T., Georges V. Houngbonon, Marc Ivaldi, and Paul T. Scott, “Market Structure, Investment, and Technical Efficiencies in Mobile Telecommunications,” In press, *Journal of Political Economy* 2024.

Fan, Ying and Chenyu Yang, “Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety,” In press, *Journal of Policitcal Economy* June 2024.

Fanou, Rodérick, Bradley Huffaker, Ricky Mok, and KC Claffy, “Unintended Consequences: Effects of Submarine Cable Deployment on Internet Routing,” in Anna Sperotto, Alberto Dainotti, and Burkhard Stiller, eds., *Passive and Active Measurement*, Springer International Publishing 2020, pp. 211–227.

Fréchette, Guillaume R., Alessandro Lizzeri, and Tobias Salz, “Frictions in a Competitive, Regulated Market: Evidence from Taxis,” *American Economic Review*, 2019, *109*, 2954–2992.

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

Hjort, Jonas and Jonas Poulsen, “The Arrival of Fast Internet and Employment in Africa,” *American Economic Review*, 2019, *109*, 1032–1079.

Ho, Kate, "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, 2009, *99*, 393–430.

Holmes, Thomas J., "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 2011, *79*, 253–302.

Houde, Jean-Francois, Peter Newberry, and Katja Seim, "Nexus Tax Laws and Economies of Density in E-commerce: A Study of Amazon's Fulfillment Center Network," *Econometrica*, 2022, *In Press.*

Jeon, Jihye, "Learning and Investment under Demand Uncertainty in Container Shipping," *RAND Journal of Economics*, 2022, *53*, 226–259.

_ **and Marc Rysman**, "International Trade in Data on the Subsea Internet Cable Network," Unpublished manuscript, Boston University 2024.

Jia, Panle, "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 2008, *76*, 1263–1316.

Kline, Brendan, Ariel Pakes, and Elie Tamer, "Moment Inequalities and Partial Identification in Industrial Organization," in "Handbook of Industrial Organization," Vol. 4, Elsevier, 2021, pp. 345–431.

LaRiviere, Jacob, Aadharsh Kannan, and Shuang Wang, "Spatial Competition and Missing Data: An Application to Cloud Computing," Unpublished manuscript 2020.

Lee, Robin, "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," *American Economic Review*, 2013, *103*, 2960–3000.

Leighton, Tom, "Improving Performance on the Internet," *Communications of the ACM*, 2009, *52* (2), 42–51.

Lipman, Andrew and Troy Tanner, "Impact of National Security Review on Submarine Cables Landing in the United States," Suboptic 2010 Conference: Enabling the next generation of networks & services, 2010.

Malone, Jacob B., Aviv Nevo, and Jonathan W. Williams, "The Tragedy of the Last Mile: Economic Solutions toCongestion in Broadband Networks," Unpublished manuscript, University of North Carolina 2021.

Molinari, Francesca, "Microeconometrics of Partial Identification," in Steven N. Durlauf, Lars Peter Hansen, James J. Heckman, and Rosa L. Matzkin, eds., *Handbook of Econometrics*, Vol. 7, Part A 2020.

Norton, William B., *The Internet Peering Playbook: Connecting to the Core of the Internet The Internet Peering Playbook: Connecting to the Core of the Internet*, DrPeering Press, 2014.

Pakes, A., J. Porter, Kate Ho, and Joy Ishii, "Moment Inequalities and Their Application: Moment Inequalities and Their Application," *Econometrica*, January 2015, *83* (1), 315–334.

Pakes, Ariel, "Alternative models for moment inequalities," *Econometrica*, 2010, *78* (6), 1783–1822.

Rysman, Marc, "Empirics of Business Data Services," Appendix B of *Business Data Services Federal Notice of Proposed Rulemaking*, FCC 16-54 2016.

Smith, Jonathan M., "Fighting Physics: A Tough Battle," *Communications of the ACM*, 2009, *52* (7), 60–65.

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

Online Appendix for "Investment and Usage of the Subsea Internet Cable Network"

Jihye Jeon and Marc Rysman

April 2025

A Additional Figures and Tables

Figure A.1: Subsea cable investment and bandwidth usage

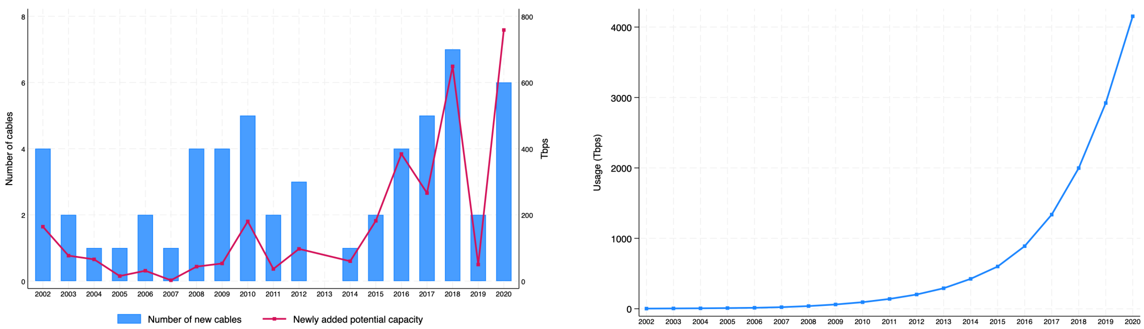

(a) Total investment (b) Total international used bandwidth

Notes: Panel (a) shows the total number and capacity of new cables by year. Panel (b) shows the trend in total international used bandwidth.

Figure A.2: Profits to a de novo and a local monopolist and welfare changes from new cable entry

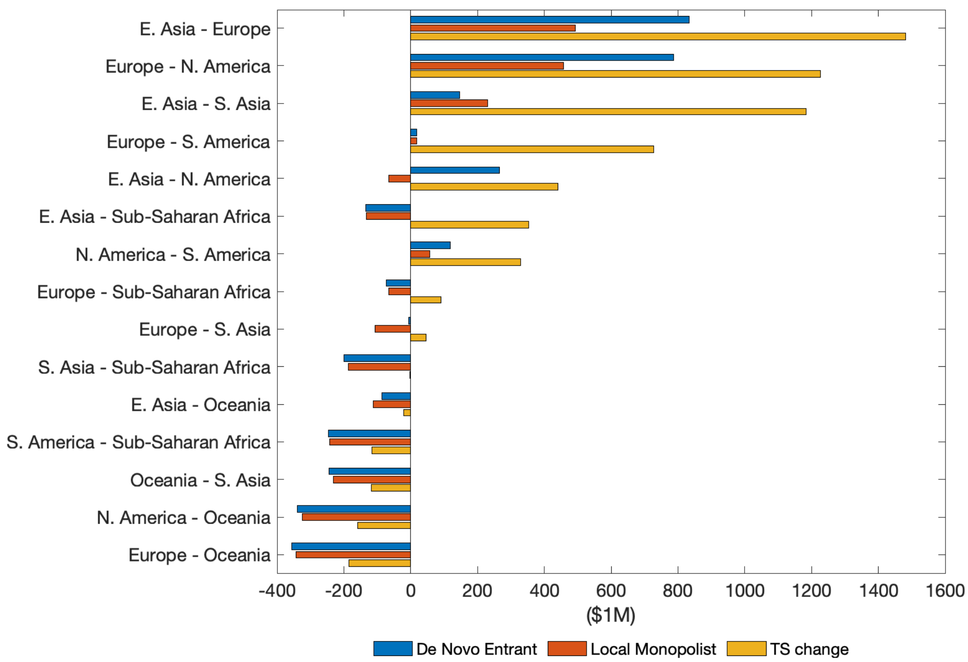

Notes: We consider an entry of a new cable in each market separately. For each market, the first bar represents the profit that a de novo entrant would gain from the new cable construction; the second bar represents the profit to a local monopolist who owns all existing cables in the market and no other cable outside the market; and the last bar the change in total surplus. The markets are ordered by the total surplus change. All figures are based on the lower bounds of the estimates and are in millions of US dollars.

Table A.1: Path decomposition

Panel A: Decomposition by endpoint regions for used bandwidth in Europe - North America in 2020

Market	Origin region	Destination region	Share (%)
Europe - North America	Europe	North America	85.92
Europe - North America	East Asia	North America	4.13
Europe - North America	East Asia	Europe	3.21
Europe - North America	Europe	Japan	1.95
Europe - North America	Europe	South Korea	0.66
Europe - North America	Japan	North America	0.64
Europe - North America	North America	Singapore	0.63
Europe - North America	North America	South Korea	0.55
Europe - North America	North America	Taiwan	0.51
Europe - North America	East Asia	South America	0.49

Panel B: Decomposition by path for data flows in Europe - South America in 2020

Path	Share (%)
Europe - North America - - South America	90.73
Europe - East Asia - North America - South America	5.43
Europe - Sub-Saharan Africa - - South America	1.92
Europe - North America - Singapore - South America	0.66
Europe - North America - South Korea - South America	0.38
Europe - Japan - North America - South America	0.28
Europe - North America - Taiwan - South America	0.21
Europe - North America - Philippines - South America	0.21
Europe - Australia - North America - South America	0.17
Europe - South Asia - Sub-Saharan Africa - South America	0.02

Notes: The table in Panel A decomposes bandwidth usage on cables connecting North America and Europe to data flows for different endpoint pairs. The table in Panel B shows the share of total data flows between Europe and South America traveling on different cable paths.

Table A.2: Capacity and cost regressions

	ln(capacity)	ln(cost)
Time	0.144***	-0.028***
	(0.027)	(0.011)
ln(capacity)		0.137***
		(0.044)
ln(length)		0.941***
		(0.045)
Constant	7.764***	9.845***
	(0.993)	(0.520)
Region fixed effects	Y	N
Observations	71	68
R-squared	0.553	0.847

Notes: The sample for each regression includes market-years with usable data and observed investment in single market cables. The capacity regression includes region fixed effects. Robust standard errors are in parenthesis. Stars indicate p-values: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.3: Internet transit price

Year	Internet transit price ($/Mbps)	% decline
1998	1,200.00	
1999	800.00	33%
2000	675.00	16%
2001	400.00	41%
2002	200.00	50%
2003	120.00	40%
2004	90.00	25%
2005	75.00	17%
2006	50.00	33%
2007	25.00	50%
2008	12.00	52%
2009	9.00	25%
2010	5.00	44%
2011	3.25	35%
2012	2.34	28%
2013	1.57	33%
2014	0.94	40%
2015	0.63	33%

Source: The Internet Peering Playbook: Connecting to the Core of the Internet, 2014, William B. Norton.

Table A.4: Alternative specifications of the supply model

Panel A: Assumptions about structural errors

	Carriers	Content providers
2002-2009	[2798.6 , 3510.2]	NA
2010-2017	[34.0 , 88.8]	[24.2 , 86.5]
2018-2020	[12.9 , 14.4]	[5.3 , 14.4]

Panel B: Alternative choice for differencing inequalities

	Nearest 1		Nearest 10	
	Carriers	Content providers	Carriers	Content providers
2002-2009	[2738.4 , 3517.9]	NA	[2674.3 , 3517.9]	NA
2010-2017	[18.8 , 98.8]	[25.2 , 121.3]	[22.3 , 98.8]	[11.3 , 96.0]
2018-2020	[12.0 , 15.6]	[3.0 , 14.9]	[12.0 , 15.6]	[5.0 , 18.8]

Panel C: Reweighted inequalities

	Carriers	Content providers
2002-2009	[2528.8 , 3261.0]	NA
2010-2017	[26.2 , 89.3]	[13.7 , 86.9]
2018-2020	[12.0 , 14.6]	[5.2 , 17.7]

Panel D: Major carriers and small carriers

	Major carriers	Small carriers
2002-2009	[2344.7 , 3633.7]	[1919.3 , 3531.3]
2010-2017	[23.4 , 86.2]	[46.9 , 89.0]
2018-2020	[6.1 , 14.0]	[13.6 , 14.4]

Panel E: Endogenous capacity

	Carriers	Content providers
2002-2009	[1811.4 , 3677.2]	NA
2010-2017	[19.3 , 100.9]	[26.5 , 98.0]
2018-2020	[13.6 , 17.1]	[3.6 , 16.4]

Notes: This table presents results for various alternative specifications for our supply model. For each specification, we report the 95% confidence interval for the markup parameter, γ_{jt}, capturing the 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. The estimates are missing for the 2002-2009 period for content providers due to the lack of observed investment episodes. The sample includes 56,935 firm-cable-market-year-level observations for carriers and 1,180 for content providers. In Panel A, we assume that $\Delta\nu_{jmt} = 0$. In Panel B, we vary the number of observations of $j'mt$ for each jmt that we choose to form the inequalities in Equation (18). In Panel C, we reweight the observations so that each market is represented equally. In Panel D, we divide cable owners into two groups: 27 major carriers including 4 content providers and the rest, and estimate γ_{jt} separately for these two groups. In Panel E, we assume that adding or subtracting an investor raises or lowers the capacity of the new cable by 5%. For analysis in Panels D-E, we make the assumption that $\Delta\nu_{jmt} = 0$ as in Panel A.

Table A.5: Counterfactuals: New cable entry by market in 2005

Market	(1) Entrant profit	(2) Change in incumbent profit: Investment market	(3) Change in incumbent profit: RoW
Overall	[-230.0, -200.5]	[49.9, 65.7]	[101.1, 133.0]
Europe - N. America	[270.9, 459.2]	[548.7, 721.8]	[53.6, 70.5]
E. Asia - Europe	[-71.5, 15.6]	[15.8, 20.8]	[533.7, 702.0]
N. America - S. America	[-145.2, -135.1]	[32.6, 42.9]	[1.4, 1.8]
E. Asia - S. Asia	[-162.2, -162.2]	[-0.0, -0.0]	[0.0, 0.1]
Europe - S. Asia	[-175.0, -156.9]	[79.9, 105.1]	[394.6, 519.1]
Europe - Sub-Saharan Africa	[-219.0, -217.8]	[-1.3, -1.7]	[0.2, 0.3]
E. Asia - Oceania	[-221.2, -217.3]	[4.0, 5.2]	[6.4, 8.4]
S. Asia - Sub-Saharan Africa	[-273.2, -272.8]	[0.2, 0.3]	[4.7, 6.1]
Europe - S. America	[-286.0, -252.9]	[0.7, 0.9]	[254.1, 334.2]
E. Asia - N. America	[-295.5, -217.6]	[50.8, 66.8]	[24.0, 31.6]
Oceania - S. Asia	[-312.8, -312.3]	[5.3, 7.0]	[19.1, 25.1]
S. America - Sub-Saharan Africa	[-314.5, -309.9]	[0.2, 0.2]	[43.6, 57.3]
E. Asia - Sub-Saharan Africa	[-350.2, -338.5]	[1.4, 1.8]	[113.2, 148.9]
N. America - Oceania	[-435.5, -431.8]	[1.6, 2.0]	[28.0, 36.8]
Europe - Oceania	[-458.4, -456.8]	[9.2, 12.2]	[40.6, 53.4]

Notes: This table reports results from a counterfactual simulation in which a de novo entrant constructs a new cable in a market at the end of 2005. Column (1) reports the entrant profit. Columns (2) and (3) report changes in the sum of incumbent profits in the market in which the investment happens and the rest of the markets, respectively. For each quantity, we report a range where the lower end is computed based on the lower bound of γ_{jt}, and the upper end is computed based on the upper bound of γ_{jt}. The first row reports the average effects over markets. The rest of the rows are ordered by entrant profit. All figures are reported in millions of US dollars.

Table A.6: Counterfactuals: New cable entry by market in 2010

Market	(1) Entrant profit	(2) Change in incumbent profit Investment market	(3) RoW
Overall	[-297.7, -256.7]	[0.9, 4.0]	[10.5, 46.8]
E. Asia - S. Asia	[-146.3, -115.1]	[-7.2, -31.9]	[4.9, 22.1]
N. America - S. America	[-163.7, -143.4]	[3.5, 15.5]	[0.4, 2.0]
Europe - S. Asia	[-208.2, -159.2]	[4.2, 18.6]	[49.2, 219.7]
Europe - Sub-Saharan Africa	[-212.4, -209.5]	[0.4, 1.8]	[0.1, 0.2]
E. Asia - Oceania	[-221.6, -214.8]	[0.1, 0.4]	[0.3, 1.4]
Europe - N. America	[-257.2, -68.7]	[14.4, 63.7]	[7.4, 33.2]
S. Asia - Sub-Saharan Africa	[-262.8, -262.5]	[0.2, 0.9]	[0.8, 3.6]
Oceania - S. Asia	[-300.6, -299.4]	[0.6, 2.5]	[2.4, 10.5]
E. Asia - Europe	[-301.6, -193.8]	[-4.4, -19.7]	[42.0, 188.6]
S. America - Sub-Saharan Africa	[-313.5, -308.5]	[-0.0, -0.0]	[4.1, 18.1]
Europe - S. America	[-361.1, -315.2]	[0.0, 0.2]	[30.1, 134.2]
E. Asia - Sub-Saharan Africa	[-366.7, -352.3]	[0.1, 0.4]	[11.5, 51.2]
N. America - Oceania	[-426.6, -422.0]	[0.1, 0.3]	[2.6, 11.7]
Europe - Oceania	[-443.1, -440.1]	[0.9, 3.8]	[4.5, 20.2]
E. Asia - N. America	[-480.3, -346.6]	[0.6, 2.7]	[-3.4, -15.1]

Notes: This table reports results from a counterfactual simulation in which a de novo entrant constructs a new cable in a market at the end of 2010. Column (1) reports the entrant profit. Columns (2) and (3) report changes in the sum of incumbent profits in the market in which the investment happens and the rest of the markets, respectively. For each quantity, we report a range where the lower end is computed based on the lower bound of γ_{jt}, and the upper end is computed based on the upper bound of γ_{jt}. The first row reports the average effects over markets. The rest of the rows are ordered by entrant profit. All figures are reported in millions of US dollars.

Table A.7: Counterfactuals: New cable entry by market in 2015

Market	(1) Entrant profit	(2) Change in incumbent profit Investment market	(3) RoW	(4) Change in consumer surplus Investment market	(5) RoW
Overall	[-171.9, 255.2]	[-20.4, -91.7]	[78.3, 349.6]	[60.2, 318.3]	[36.6, 223.9]
Europe - N. America	[236.4, 2072.7]	[-35.3, -156.2]	[16.2, 72.1]	[510.4, 2749.1]	[21.2, 251.1]
E. Asia - Europe	[21.2, 1187.4]	[-72.7, -322.1]	[298.7, 1342.7]	[44.0, 226.4]	[156.3, 892.0]
E. Asia - S. Asia	[10.5, 556.5]	[-25.7, -113.9]	[327.2, 1459.4]	[0.1, 0.5]	[107.9, 600.4]
N. America - S. America	[-74.3, 227.9]	[20.2, 89.7]	[2.3, 10.6]	[110.5, 564.7]	[5.2, 79.6]
Europe - S. Asia	[-112.6, 231.6]	[-79.9, -353.8]	[34.3, 152.7]	[6.3, 31.7]	[12.3, 71.6]
E. Asia - N. America	[-163.2, 982.6]	[-131.6, -601.2]	[-78.5, -347.7]	[183.9, 959.0]	[26.7, 217.4]
E. Asia - Oceania	[-187.4, -96.4]	[-1.5, -6.5]	[0.4, 1.9]	[25.1, 126.1]	[1.3, 20.6]
Europe - Sub-Saharan Africa	[-188.3, -133.9]	[6.9, 30.4]	[0.7, 3.2]	[23.2, 116.3]	[1.0, 18.1]
Europe - S. America	[-199.8, 344.5]	[0.2, 1.1]	[292.9, 1301.4]	[0.0, 0.0]	[120.3, 670.0]
S. Asia - Sub-Saharan Africa	[-248.4, -237.1]	[2.7, 12.1]	[15.8, 70.0]	[0.1, 0.5]	[6.3, 35.0]
Oceania - S. Asia	[-280.5, -254.3]	[6.1, 27.2]	[30.7, 136.4]	[0.0, 0.0]	[11.6, 64.1]
S. America - Sub-Saharan Africa	[-281.0, -210.7]	[-0.2, -0.8]	[53.2, 235.7]	[0.0, 0.1]	[16.1, 89.4]
E. Asia - Sub-Saharan Africa	[-303.5, -126.7]	[0.6, 2.8]	[116.9, 521.6]	[0.0, 0.0]	[38.8, 215.7]
N. America - Oceania	[-396.7, -351.8]	[-3.1, -13.7]	[16.2, 72.4]	[0.0, 0.0]	[8.0, 44.6]
Europe - Oceania	[-411.4, -364.4]	[6.8, 30.1]	[47.5, 211.6]	[0.0, 0.0]	[16.1, 89.4]

Notes: This table reports results from a counterfactual simulation in which a de novo entrant constructs a new cable in a market at the end of 2015. Column (1) reports the entrant profit. Columns (2) and (3) report changes in the sum of incumbent profits in the market in which the investment happens and the rest of the markets, respectively. Columns (4) and (5) report changes in consumer surplus in the market in which the investment happens and the rest of the markets, respectively. For each quantity, we report a range where the lower end is computed based on the lower bound of γ_{jt}, and the upper end is computed based on the upper bound of γ_{jt}. The first row reports the average effects over markets. The rest of the rows are ordered by entrant profit. All figures are reported in millions of US dollars.

Table A.8: Counterfactuals: New cable entry by market in 2020

Market	(1) Entrant profit	(2) Change in incumbent profit	(3)	(4) Change in consumer surplus	(5)
		Investment market	RoW	Investment market	RoW
Overall	[32.1, 126.9]	[-69.1, -107.8]	[176.3, 257.6]	[114.9, 154.4]	[97.6, 135.4]
E. Asia - Europe	[833.1, 1174.6]	[-339.9, -441.9]	[328.5, 594.9]	[415.2, 556.9]	[244.4, 343.5]
Europe - N. America	[786.7, 1108.5]	[-328.7, -618.2]	[54.4, 74.1]	[631.3, 853.8]	[82.1, 120.1]
E. Asia - N. America	[266.2, 488.9]	[-331.2, -497.5]	[102.5, 176.4]	[163.3, 220.0]	[240.5, 331.0]
E. Asia - S. Asia	[146.6, 233.2]	[83.3, 108.2]	[664.0, 912.6]	[1.4, 1.8]	[288.2, 393.2]
N. America - S. America	[117.7, 199.7]	[-60.2, -88.2]	[31.3, 45.3]	[200.7, 268.8]	[38.8, 57.9]
Europe - S. America	[17.2, 125.3]	[0.3, 0.5]	[519.6, 760.7]	[0.0, 0.0]	[190.0, 259.2]
Europe - S. Asia	[-6.2, 53.1]	[-101.2, -131.6]	[47.6, 62.5]	[75.4, 100.7]	[30.8, 44.2]
Europe - Sub-Saharan Africa	[-73.1, -36.4]	[6.7, 8.7]	[6.6, 9.2]	[136.8, 182.7]	[13.5, 22.3]
E. Asia - Oceania	[-86.9, -51.5]	[-24.7, -33.0]	[-16.0, -23.4]	[97.9, 130.7]	[8.8, 14.7]
E. Asia - Sub-Saharan Africa	[-134.2, -72.4]	[1.5, 1.9]	[360.1, 497.8]	[0.0, 0.0]	[126.7, 172.8]
S. Asia - Sub-Saharan Africa	[-200.5, -188.3]	[14.3, 18.6]	[127.1, 163.7]	[0.6, 0.8]	[56.4, 76.9]
Oceania - S. Asia	[-243.7, -234.0]	[12.6, 16.3]	[82.6, 107.9]	[0.0, 0.0]	[30.3, 41.3]
S. America - Sub-Saharan Africa	[-246.6, -234.0]	[4.0, 5.2]	[99.8, 144.5]	[0.2, 0.3]	[25.8, 35.2]
N. America - Oceania	[-338.9, -322.9]	[14.1, 18.3]	[117.5, 169.2]	[0.0, 0.0]	[47.6, 64.9]
Europe - Oceania	[-356.2, -341.0]	[12.6, 16.4]	[119.4, 168.1]	[0.0, 0.0]	[39.3, 53.6]

Notes: This table reports results from a counterfactual simulation in which a de novo entrant constructs a new cable in a market at the end of 2020. Column (1) reports the entrant profit. Columns (2) and (3) report changes in the sum of incumbent profits in the market in which the investment happens and the rest of the markets, respectively. Columns (4) and (5) report changes in consumer surplus in the market in which the investment happens and the rest of the markets, respectively. For each quantity, we report a range where the lower end is computed based on the lower bound of γ_{jt}, and the upper end is computed based on the upper bound of γ_{jt}. The first row reports the average effects over markets. The rest of the rows are ordered by entrant profit. All figures are reported in millions of US dollars.

B Constructing Path Characteristics

In this section, we describe how we construct path characteristics. We first describe the construction of the capacity measure. 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 the capacity for the cables in that path (that is, the capacity for the cables in that market, because 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 take the minimum capacity of market-level capacities of each market on a path as the path capacity.[25]. We explore alternative measures such as the average

[25]For example, if we are considering the path between A and C that goes through B, and the potential capacity on AB is 1 and BC is 10, we take the potential capacity of the ABC path to be 1. Thus, low-capacity markets create bottlenecks. This use of the minimum distinguishes our approach

and weighted average (by length) in Jeon and Rysman (2024). We construct other measures of cable features in a similar way. We construct the distance of a path as the sum of region pair distances where the region pair distance is computed as the average distance between countries in those regions. 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 summarize these variables for 3,646 paths in our data in Table B.9.

Table B.9: Summary of path-level data

	Mean	SD	5th	95th	N
Capacity (in Tbps)	28.69	57.28	0	150	41,012
Length (in km)	21,704.67	8,685.37	6,840	34,933	41,012
Number of cables	6.35	3.97	1	14	41,012

Notes: This table summarizes paths in our sample. The observations are at the path-year level. There are 3,646 paths in total.

C Model with Subregions

For readability, we present the model at the level of the region-pair, or what we term the market. However, in order to make use of island countries, we solve actually estimate the demand model at the level of the subregion, as described in Section 4.1. In this section, we present the mathematics of the demand model in terms of subregions. We reuse notation to some extent.

We group countries into *subregions* $g = 1, \dots, G$ and subregions into *regions* $r = 1, \dots, R$. The subregions form a partition of the countries, and the regions form a partition of the subregions. As described above, 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.

The definition of $d_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ is the same as above. Paths are a set of subregion pairs for which there is an active cable. We restrict ourselves to paths that go through

from the multiplicative iceberg costs of Allen and Arkolakis (2022) four subregions at most. We define $\mathcal{L}_{pckt}$ to be the set of subregion pairs on path p between country pair ck. Path attractiveness $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ depends on the level of capacity on the path, $K^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$. We assume $K^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \min_{l \in \mathcal{L}_{pckt}} K_{lt}(\boldsymbol{a}_{lt}, K_{lt-1})$, now defined over subregion pairs rather than region pairs. We have an analogous adjustment for $n^{\delta}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$. The term $\delta_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$ is otherwise defined as before, as is the share of data going from c to k on path p at time t $s_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$, the quality of connection $v_{ckt}(\boldsymbol{A}_t, \mathcal{S}_t)$, and data traveling on a path between two countries $\hat{d}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t)$

The total demand between countries c and k traveling on subregion pair l in period t is now obtained by:

$$\tilde{d}_{lckt}(\boldsymbol{A}_t, \mathcal{S}_t) = \sum_{p=1}^{P_{ckt}} \hat{d}_{pckt}(\boldsymbol{A}_t, \mathcal{S}_t) \mathbb{1}\{l \in \mathcal{L}_{pckt}\}.$$

Let l index pairs of subregions, and let $\mathcal{R}_m$ 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*, $\mathcal{R}_m$ contains three subregion pairs, *Japan-North America*, *Korea-North America* and *Rest of East Asia-North America*.

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

$$D_{mt}(\boldsymbol{A}_t, \mathcal{S}_t) = \sum_{l \in \mathcal{R}_m} \left(\sum_{c=1}^{C-1} \sum_{k=c+1}^{C} \tilde{d}_{lckt}(\boldsymbol{A}_t, \mathcal{S}_t) \right). \tag{C.19}$$

D Inference Procedure

We implement the conditional chi-squared (CC) test proposed by Cox and Shi (2023) to compute confidence sets for the true parameter value μ^* based on a set of unconditional moment inequalities $l = 1, ..., L$. Moment inequalities are described in Section 4.3. 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

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

We denote $\hat{\Sigma}(\gamma)$ as an estimator of the 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 $\gamma \in \mathcal{A}$. 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:

$$T(\gamma) = \min_{\mu:\mu \geq 0} N(\overline{m}(\gamma) - \mu)'\hat{\Sigma}(\gamma)^{-1}(\overline{m}(\gamma) - \mu) \tag{D.20}$$

where N is the sample size. The vector μ is of dimensions $L \times 1$. When defining the objective function of the minimization problem in Equation (D.20), it occasionally happens that the covariance matrix $\hat{\Sigma}(\gamma)$ is singular or close to singular. When this is the case, we follow Andrews and Barwick (2012), and substitute the covariance matrix in Equation (D.20) for the following matrix:

$$\tilde{\Sigma}(\gamma) = \hat{\Sigma}(\gamma) + \max\{0.012 - \det(\hat{\Omega}(\gamma), 0\} Diag(\hat{\Sigma}(\gamma)) \tag{D.21}$$

where

$$\hat{\Omega}(\gamma) = Diag^{-\frac{1}{2}}(\hat{\Sigma}(\gamma))\hat{\Sigma}(\gamma)Diag^{-\frac{1}{2}}(\hat{\Sigma}(\gamma)),$$

$Diag^{-\frac{1}{2}}(\hat{\Sigma}(\gamma))$ is a matrix such that $Diag^{-\frac{1}{2}}(\hat{\Sigma}(\gamma))Diag^{-\frac{1}{2}}(\hat{\Sigma}(\gamma)) = Diag^{-1}(\hat{\Sigma}(\gamma))$, and $Diag(\hat{\Sigma}(\gamma))$ is the $L \times L$ diagonal matrix whose diagonal elements are equal to those of $\hat{\Sigma}(\gamma)$.

Step 4: Count how many values of μ equal 0. We denote this number as $\hat{r}$.

Step 5: Accept/reject γ. 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 $\mathcal{A}$.

E Consumer Welfare

This section describes the back-of-the-envelope calculation for the change in consumer surplus from constructing new cables. From the Lerner Index, we know that

$$\frac{P - MC}{P} = -\frac{1}{\epsilon} \tag{E.22}$$

where P is price, MC is marginal cost, and ϵ is the own-price elasticity, which we take to be $\epsilon > 1$. In our supply model, we estimate γ, which we interpret as $P - MC$. Thus, to the extent that we know the price level P, we can infer the elasticity ϵ from the Lerner Index and γ.

Under the assumption of constant elasticity of demand, we can further calculate welfare. Suppose the demand function is $Q = AP^{-\epsilon}$, so that $\frac{\partial Q}{\partial P}\frac{P}{Q} = -\epsilon$. Then, consumer welfare W is:

$$W = \int_P^\infty Q(s)ds = \int_P^\infty As^{-\epsilon}ds = \frac{As^{1-\epsilon}}{1-\epsilon}\Big|_{s=P}^{\infty} + C = \frac{PQ}{\epsilon - 1} + C,$$

where C is an integrating constant. Thus, if W_i is consumer welfare under policy i, then the change in consumer welfare

$$\Delta W = W_2 - W_1 = \frac{P_2Q_2 - P_1Q_1}{\epsilon - 1}. \tag{E.23}$$

To compute consumer surplus based on the data we have available, we first construct a price index. TeleGeography provides a sample of median monthly lease prices on a sample of routes from 2015 to 2019 for two types of products: 10 Gbps wavelength and 100 Gbps wavelength. We take the average over all relevant interregional routes (18 routes) to construct the yearly average price per capacity, P. We assume that the 2020 prices are equal to the levels in 2019. Then, we compute ϵ based on Equation (E.22) and our estimates of γ for carriers. We approximate the price under a new policy (P_2) using the definition of elasticity as

$$P_2 = \frac{(Q_2 - Q_1)P_1}{Q_1\epsilon} + P_1$$

where P_1 and Q_1 are observed price and quantity in the data and Q_2 is the predicted quantity under the new policy. Then, we compute the change in consumer welfare using Equation (E.23).

F Dynamic Model

There are several challenges in estimating a dynamic model for our setting. Since markets are interdependent, the state includes relevant information about all markets. Specifically, it includes firm-level capacities for all firms and the number of cables in all markets. Moreover, the problem is highly non-stationary due to the changing landscape of the industry. The size of cables and the demand for data exchanges are growing rapidly, while the unit cost of laying a cable is falling in our sample period. Content providers enter the market starting in 2010 and grow to be major investors by the end of our sample. All of these factors make it infeasible to rely on a two-step estimation method (e.g. Bajari, Benkard and Levin (2007)), which requires a consistent estimator of firms' investment policy as a function of the state variables.

We make a few relatively strong assumptions to incorporate dynamics. First, we assume that firms play a game of full commitment. That is, at the beginning of the game, firms choose all of their actions in each time period. Second, we assume firms have perfect foresight over the evolution of demand and cost so firms know the evolution of demand and costs. Note that firms still observe profits with measurement error ϵ_{ajmt}. Third, we assume that after some period T, firms receive the same profit in all subsequent periods.

The game begins with an allocation of capacities and cables, so the state space at the beginning of the game is $\mathcal{S}_0 = \{\{k_{jm0}\}_{j=1,\ldots,J,\, m=1,\ldots,M},\, \{n_{m0}\}_{m=1,\ldots,M}\}$. The profit to firm j at the start of the game is:

$$\Pi_j\left(\{\boldsymbol{A}_t\}_{t=1}^T, \mathcal{S}_0\right) = E\left[\sum_{t=1}^{\infty} \beta^t \bar{\pi}_{jt}\left(\boldsymbol{A}_t, \mathcal{S}_t\right) \Big| \mathcal{J}_{j0}\right],$$

where where $t > T$ implies $\boldsymbol{A}_t = \boldsymbol{A}_T$ and $\mathcal{S}_t = \mathcal{S}_T$. Before T, $\mathcal{S}_t$ evolves according to Equations (1), (2), and (4) as above. The term $\mathcal{J}_{j0}$ is the information set j holds at the start of the game, which includes future determinants of demand and cost (except for rival strategies).

Firms play a Nash equilibrium, simultaneously solving:

$$\max_{\{a_{jmt}\}_{m=1,\ldots,M;\ t=1,\ldots,T}} \Pi_j\left(\{\boldsymbol{A}_t\}_{t=1}^T, \mathcal{S}_0\right)$$

Nash equilibrium implies that for firm j that chooses to invest in period τ in market m the investment is more profitable than not investing:

$$\Pi_j\left(\left\{\{\boldsymbol{A}_t\}_{t=1}^{\tau-1}, \boldsymbol{A}_\tau^{j,m,1}, \{\boldsymbol{A}_t\}_{t=\tau+1}^{T}\right\}, \mathcal{S}_0\right) - \Pi_j\left(\left\{\{\boldsymbol{A}_t\}_{t=1}^{\tau-1}, \boldsymbol{A}_\tau^{j,m,0}, \{\boldsymbol{A}_t\}_{t=\tau+1}^{T}\right\}, \mathcal{S}_0\right) \geq 0$$

Observations of no investment lead to an analogous reversed inequality. Note that a deviation in a given period leads to a change of state variables in all following periods and so affects the whole stream of profits in following periods. However, we do not consider how rivals would change future decisions in response because of the nature of Nash equilibrium arising from our assumption of commitment.

We use the same differencing strategy as expressed in Equation (14) to address selection in the dynamic context. The moment that we bring to data (the analog to Equation (14) under dynamics) is:

$$\begin{aligned} E\Bigg[&\sum_{z\in\mathcal{M}} \Delta r_{jmzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - \sum_{z\in\mathcal{M}} \Delta r_{j'mzt}(\boldsymbol{A}_t, \mathcal{S}_t; \gamma_{jt}) - c_{jmt}\left(\boldsymbol{a}_{mt}^{j,1}\right) + c_{j'mt}\left(\boldsymbol{a}_{mt}^{j',1}\right) + \\ &\sum_{\tau=t+1}^{\infty} \beta^{\tau-t} \sum_{m\in\mathcal{M}} \left(r_{jm\tau}(\boldsymbol{A}_\tau, \mathcal{S}_\tau; \gamma_{jt}) - r_{jm\tau}(\boldsymbol{A}_\tau, \tilde{\mathcal{S}}_\tau; \gamma_{jt})\right) - \left(r_{j'm\tau}(\boldsymbol{A}_\tau, \mathcal{S}_\tau; \gamma_{jt}) - r_{j'm\tau}(\boldsymbol{A}_\tau, \tilde{\mathcal{S}}'_\tau; \gamma_{jt})\right) \\ &\Bigg| a_{jmt} = 1, a_{j'mt} = 0, \mathcal{Z}_{jt}, \mathcal{Z}_{j't} \Bigg] \geq 0. \end{aligned} \quad \text{(F.24)}$$

In this equation, the first and third lines are the same as in Equation (14). The second line is new and captures the future stream of revenues arising from the comparison. The first parenthesis captures the change in revenue to j from its investment in market m and period t. The second parenthesis captures the change in the revenue of j' from its deviation. We use $\tilde{\mathcal{S}}_\tau$ to denote the state variables that follow j's deviation in period t and $\tilde{\mathcal{S}}'_\tau$ for the state variables that follow the deviation of j'.

Equation (F.24) is the moment inequality we bring to data. Each case of two firms making different choices in the same market and time period potentially provides an observation. We use the five nearest neighbors in terms of changes in quantity in the current period, as described for our static model. In practice, we apply the moment separately for each of our three major time periods. That is, for the time period

2002-2009, we treat the game as starting in 2002 with T representing 2009 and the firm believing that state variables will remain the same after 2009. Then we repeat the procedure for the time period 2010-2017 and again for 2018-2020. We believe these are reasonable periods over which firms formed consistent strategies. We also estimate the model separately for content providers and non-content providers.

Note that within these groups, γ_{jt} is constant, so it is correct in Equation F.24 to use a single parameter, such as using γ_{jt} in the revenue function of firm j'. We prefer to write this way to emphasize that we estimate a single parameter at a time. We assume that firms believe γ_{jt} will stay the same after period T. An alternative approach is to allow γ_{jt} and data flows to evolve after period T. However, in our results, they are largely offsetting, with γ_{jt} falling about 30% per year and data flows growing by about 30% year. Thus, assuming they are constant after T is a reasonable approximation.

We set $\beta = 0.95$ and estimate bounds using the same procedure described in Section 4.3. Panel A of Table F.10 presents the results. We conduct a robustness check in which we drop observations from the final year in each time period in the estimation. This is to ensure that our results are not driven by the assumption that the profits will stay constant after the final year. Panel B reports the confidence intervals, which are robust to this sample change for the first two periods. The bounds are higher for the 2018-2020 period, which could be due to the markup falling more rapidly in this period or the smaller number of years included in the sample.

Table F.10: Dynamic model estimates

Panel A: Full sample

	Carriers	Content providers
2002-2009	[32.5 , 76.6]	NA
2010-2017	[0.6 , 2.4]	[0.3 , 2.4]
2018-2020	[0.5 , 0.7]	[0.3 , 0.9]

Panel B: Omitting observations from final year

	Carriers	Content providers
2002-2009	[21.6 , 77.8]	NA
2010-2017	[0.7 , 2.6]	[0.7 , 2.9]
2018-2020	[0.6 , 0.7]	[0.5 , 0.9]

Notes: We report the 95% confidence set for the markup parameter, γ_{jt}, capturing the profit (measured in millions of US dollars) from an increase in the quantity served by the firm by one Tbps. The estimates are missing for the 2002-2009 period for content providers due to the lack of observed investment episodes. The sample includes 56,935 firm-cable-market-year-level observations for carriers and 1,180 for content providers.