Diversification, Market Entry, and the Global Internet Backbone

El Hadi Caoui — U of Toronto

Andrew Steck — U of Toronto

January 14, 2025*

Abstract

In many industries, buyers diversify their supplier base to manage supplier disruption risk. We investigate the importance of such diversification as a determinant of demand and supplier entry in the context of the internet backbone, the worldwide network of undersea fiber-optic cables that underpins the internet. We specify a model of international bandwidth demand and cable operators' dynamic entry and supply choices. The model is estimated using novel data on cross-border data flows, prices, cable characteristics, and disruptions. Counterfactual analysis reveals that supplier diversification accounts for a large share of entry and surplus created between 2005 and 2021. Relative to the socially optimal level of entry, distortions due to suppliers' inability to capture fully the social benefits of diversity are as large as distortions due to business stealing.

Keywords: diversification, supply disruptions, internet backbone, entry, dynamic games

JEL Classification: L13, L96, L86

*Caoui: elhadi.caoui@rotman.utoronto.ca. Steck: andrew.steck@rotman.utoronto.ca. Department of Management, University of Toronto Mississauga and Rotman School of Management. 105 St George St, Toronto, ON M5S 3E6. We are grateful to Tim Bresnahan and Mitsuru Igami for extensive comments. We thank Ambarish Chandra, Mert Demirer, Avi Goldfarb, Shane Greenstein, Matthew Mitchell, Julie Holland Mortimer, Andrew Sweeting, Joel Waldfogel, Xiaosong Wu, Biliana Yontcheva, seminar and conference participants at the University of Toronto, University of Maryland, NYU Stern, Cornell Johnson, University of Virginia, Bocconi, IIOC, WEAI, CEA, CEPR Applied IO Conference, BSE Summer Forum, CRESSE, NYU Stern Economics of Strategy workshop, and the NBER IO Spring Meeting 2024 for helpful comments and discussions. We are grateful to Lilian Lea and Alan Mauldin from TeleGeography for helping us access the data and answering our questions about the global undersea internet cable industry. The analysis, findings, and conclusions expressed in this paper should not be attributed to TeleGeography.

1 Introduction

Supply disruptions are a prevalent concern in many industries, where natural disasters or production failures can have economically devastating consequences. Recent examples highlighting the vulnerability of supply networks abound. In March 2021, the container ship *Ever Given* ran aground and blocked the Suez Canal for over a week; a political protest in February 2022 shut down traffic on the Ambassador Bridge, which carries 25% of trade between the U.S. and Canada; and disruptions due to COVID-19 mitigation measures and outbreaks significantly disrupted chip supply in industries from graphic cards to automobiles. In response to these disruptions, buyers commonly diversify their sourcing across multiple suppliers to mitigate risks. In this paper, we examine empirically how diversification affects suppliers' incentives to enter and compete, and more broadly, its long-run impact on market structure. We do so in an important but under-studied setting: the undersea telecommunication cables that comprise the global "backbone of the internet."

The effect of supplier diversification on market structure operates through two primary channels. First, as long as the disruptions of different suppliers are not perfectly correlated, diversification increases the degree of horizontal differentiation between suppliers; this differentiation insulates suppliers from competitive pressure. Second, for a fixed level of prices, market entry by additional suppliers will increase the aggregate quantity demanded, a "market expansion" effect; intuitively, the value of the "portfolio" of the market supply increases with the number of competitors.

Our empirical application focuses on the global internet backbone from 2005 to 2021. This worldwide network of undersea fiber-optic cables comprises the primary means of intercontinental information transport, carrying more than 98% of all international internet traffic (data, video calls, instant messages, and emails). Operational undersea cable networks are essential to well-functioning global economic and financial systems: e.g., the U.S. Clearing House Interbank Payment System processes more than \$10 trillion per day in transactions with more than 22 economies via undersea cables (Federal Communications Commission (2015)). Undersea cables also feature prominently in intra-continental and even intra-national communication networks. Figure 1 shows the global undersea cable network in 2022.[1]

The undersea internet cable industry is well-suited to examine the effect of the demand for diversification on market structure. Hundreds of cable failures and repairs take place every year; given the risk to international trade, financial markets, and social welfare, the

[1] We do not consider the network of terrestrial cables, arguably another important component of the internet backbone, due to a lack of data on terrestrial cables and the relatively more complex topology of the terrestrial cable network.

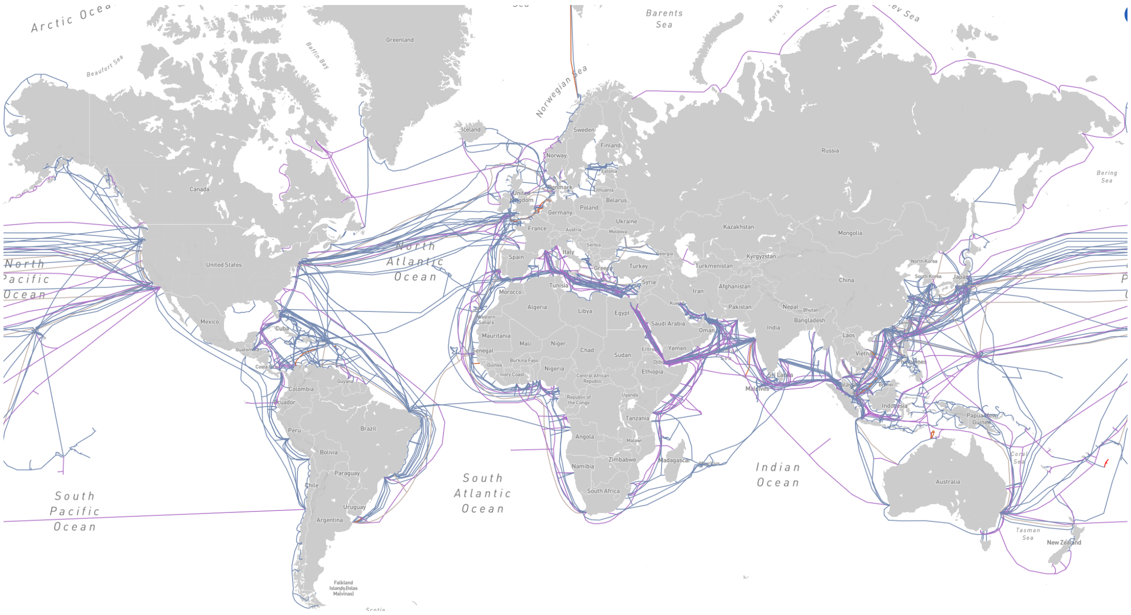

Figure 1: The Global Internet Backbone (source: Infrapedia.com)

industry therefore places a high value on network resilience and path diversity. Industry estimates put the financial impact from interruptions of undersea fiber-optic systems in excess of $1.5 million per hour (Malphrus (2009)). In addition, the industry has witnessed significant growth over the past twenty years, providing many instances of entry in a variety of market conditions. Finally, the undersea cable industry is of interest in its own right as a critical component of the "industrial organization of the internet supply chain."

We leverage comprehensive data on this industry, which were obtained from TeleGeography, a telecommunications market research firm. The data provide detailed information on undersea cables (e.g., construction costs, cable length, ready-for-service dates, landing stations, capacities), as well as extensive quarterly panels on bandwidth prices (at the city-pair level) and data flows (at the country-pair level). To the authors' knowledge, these data have not previously been used in published economics research. We supplement this main data source with an array of demand and cost factors that are either specifically relevant for our setting (e.g., broadband subscriptions, data centers, internet exchange points, electricity prices) or typically used in the international trade literature to predict trade flows (as in Blum and Goldfarb (2006)). The scope of these data on both the supply and demand side allows us to estimate a detailed and flexible bandwidth demand system and a dynamic structural model of entry and competition between undersea cable operators.

Our empirical model considers the industry through the lens of a dynamic oligopoly game, with repeated Cournot competition between cable operators taking place each period in each market. We use a non-stationary Markov Perfect equilibrium concept to accommodate the evolution of exogenous demand and cost factors in our setting: our data covers much of the industry's earliest growth and witnesses substantial increases in demand and steady decreases in costs. We define a market to be a country-pair and use details on cable landing points to match cables to markets. The richness of the data allows us to control for many dimensions of heterogeneity, including market-level unobserved heterogeneity and regional-time shocks.

Estimation proceeds in three steps. Bandwidth demand is estimated as a function of a host of market characteristics, bandwidth prices, and the number of cables serving the market; we employ cost-shifting instrumental variables to address price endogeneity. From these estimates, we recover price elasticities and buyers' preference for diversity. The demand model achieves a high degree of fit and the estimated parameters are consistent with predictions of a consumer-level utility maximization problem under costly supply disruptions. We find decreasing marginal returns from diversification: e.g., entry of a second cable (holding prices fixed) expands demand as much as a 28.3% decrease in bandwidth prices, whereas entry of the eighth cable is equivalent to a 7.5% decrease in prices.

Next, we recover the marginal costs of bandwidth implied by the demand estimates and the first-order conditions of the firms' profit maximization. Third, we incorporate these profits into the dynamic game of entry and competition between cable operators and use a nested pseudo-likelihood routine (the NPL algorithm, Aguirregabiria and Mira (2007)) to estimate the dynamic investment costs and distribution of firms' private information shocks. To address concerns around convergence issues associated with the NPL algorithm (Pesendorfer and Schmidt-Dengler (2010)), we implement several alternative estimators: e.g., two-step estimators (e.g. Pesendorfer and Schmidt-Dengler (2008)) and the spectral algorithm recently proposed by Aguirregabiria and Marcoux (2021). The latter estimator is an iterative algorithm that has the benefits of imposing equilibrium restrictions, is robust to fixed-point instability, and like the NPL algorithm, avoids the approximation of high-dimensional Jacobians. One identification challenge in this exercise stems from the nascent nature of the setting we study: we observe hundreds of instances of cable entry over our data's timeframe but only a handful of cables exiting. The lack of exit events presents us with an identification challenge which we address by leveraging the cable construction costs data to separately identify entry and fixed costs.

The estimated model is used for two counterfactual exercises. First, we examine how supplier diversification influences industry dynamics, cable investment, and surplus generation. In this counterfactual, buyers cannot diversify their supplier base: new cable entry increases competition and lowers prices but provides no benefits through diversification. We find supplier diversification significantly drives entry: without diversification, cable investment decreases by 12%. A substantial portion of new capacity investment is thus linked to buyers' diversification preferences. The net present value of total surplus per market over the sample period averages \$1.11 billion under the observed equilibrium. Supplier diversification accounts for 11% of total surplus and 27% of consumer surplus.

Next, we assess whether market forces provide excessive or insufficient levels of diversification. We compare entry levels under the observed equilibrium to the socially optimal level of entry. We disentangle two distortions: business-stealing effects lead to excessive entry, as entrants reduce incumbents' outputs, creating a gap between social and private benefits of entry. By contrast, diversity effects result in insufficient entry, as marginal entrants enhance surplus through diversity but fail to fully capture it without price discrimination (Spence (1976), Mankiw and Whinston (1986)). The dominant effect depends on the shape of the demand curve, entry costs, and post-entry competition. We disentangle these distortions by comparing the social planner's solution, which chooses the optimal dynamic entry path to maximizes total surplus, to a coordinated entry solution, maximizing producer surplus (in both cases taking post-entry competition as given). Because both distortions are eliminated in the first scenario and only business-stealing distortions are eliminated in the latter scenario, comparing the two counterfactuals allows us to quantify the two distortions' relative impact on market outcomes.

The expected number of cables at the end of the sample (2021-Q4) in the market outcome is distorted downward by diversity effects and upward by business stealing: the magnitude of diversity distortions in terms of number of entrants ranges from 54% to 125% of the business-stealing distortion. For most markets, business-stealing tends to dominate leading to moderately excessive entry. Relative to the market outcome, total surplus under the planner's solution is on average 10% higher: 53% of this welfare gap is due to diversity effects, whereas 47% is due to business-stealing. We analyze the cross-market heterogeneity in the relative sizes of these distortions and which market features (e.g., the size of entry costs, the market size, and demand growth over time) are more likely to lead to insufficient levels of supplier diversification. We conclude with a comparative statics exercise demonstrating the first-order impact of disruption costs on equilibrium market outcomes.

Our findings have significant implications for policy and market design, both within the telecommunications industry and in other industries where supplier reliability is a critical concern. Supplier diversification is a key feature in the equilibrium market structure of international and domestic energy transportation markets (fossil fuels or electricity) among others. Our analysis reveals that profit-maximizing suppliers are unlikely to fully internalize the diversification-related social benefits of entry decisions, and may enter at rates that are sub-optimal from a social welfare perspective; this suggests that there are instances where targeted entry subsidies would pass cost-benefit tests.

Relatedly, our analytical framework has important implications for regulators evaluating the competitiveness of market structures in industries where supplier diversification is a key concern. Regulators who fail to consider the additional demand generated by incremental suppliers–whose presence allows for risk diversification–may incorrectly conclude that a market is sufficiently competitive. Similarly, antitrust authorities assessing the competitive effects of proposed mergers in such industries must account for firms' private incentives to provide diversification, in order to reach accurate conclusions about welfare effects.

The rest of the paper proceeds as follows: after a literature review, Section 2 provides background information on the industry. Section 3 presents the data used in the empirical exercise. Section 4 describes the industry model. Section 5 discusses our estimation strategy and results. Section 6 presents the results of counterfactual exercises, and Section 7 concludes.

Literature Review. This paper contributes to several strands of the literature: the literature on dynamic games of entry and exit in non-standard demand settings; the literatures studying supply reliability and endogenous product variety; the economics of the internet and its infrastructure; and somewhat less directly, the macroeconomic and trade literature on supply chain disruptions and their propagation.

The first strand of literature is the empirical industrial organization literature on dynamic entry games, particularly those involving complexities in demand, production technology, or regulatory environments. For example, Collard-Wexler (2013) shows that cyclical demand significantly impacts the firm size distribution and market structure in the ready-mix concrete industry. Similarly, Kalouptsidi (2014) and Jeon (2022) examine bulk shipping, where time-to-build and demand uncertainty influence equilibrium prices and investment decisions. By contrast, our paper focuses on supply uncertainty and analyzes the role supplier diversification plays in shaping market structure and welfare outcomes. Methodologically, we build on the dynamic games estimation literature (Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2008)). Given the rapid growth in demand in this nascent industry, our treatment of non-stationarity aligns with studies of high-tech commodities such as Igami (2017).

This paper contributes to the industrial organization literature on supply reliability. Shepard (1987) and Farrell and Gallini (1988) propose models where a monopolist may encourage a "second source" of supply in order to credibly commit to future reliability or prices and expand demand. In the case of electricity markets, Crew and Kleindorfer (1978) provide early treatment of optimal reliability under demand uncertainty; Joskow and Tirole (2007) analyze optimal operating system reserves under potential network collapses. Lim and Yurukoglu (2018) analyze underinvestment in electricity distribution and reliability stemming from time inconsistency and moral hazard and how these distortions interact with a regulator's political affiliation. This literature typically focuses on a single seller (or planner), whereas we consider a case where the provision of reliability is decentralized and relies on many suppliers' private incentives to enter.

This paper is also related to the empirical literature on endogenous product choice (Draganska et al. (2009), Sweeting (2013), Eizenberg (2014), Berry et al. (2016), Wollmann (2018), Fan and Yang (2020)). This literature analyzes the impact of competition on product variety and shows that distortions similar to the ones we consider can lead to too many or too few products being offered. We highlight a distinct source for preferences for product variety, i.e., supply disruptions, which has received less attention in this literature.

Third, this paper contributes to the literature on the economics of the internet, particularly its infrastructure. Early theoretical studies analyzed competition, interconnection, peering, and antitrust issues in the terrestrial internet backbone market (e.g., Crémer et al. (2000), Besen et al. (2001), Laffont et al. (2001), Caillaud and Jullien (2003)). Greenstein (2015) provides a comprehensive analysis of the internet's commercialization. For undersea internet cables, Hjort and Poulsen (2019) find that the first undersea cable significantly boosts local employment, wages, and firm productivity in African locations. Regarding earlier communication technologies, Steinwender (2018), following Hoag (2006) and Garbade and Silber (1978), shows that early transatlantic telegraph cables linked markets more tightly, improving trade efficiency. Closest to our work, Jeon and Rysman (2024) models global data flow routing and quantifies network externalities. Our paper complements their analysis by examining supply disruptions and the distortions affecting diversity levels.

The disruptions examined in this paper relate to recent work on global supply chain disruptions and shock propagation in international trade and macroeconomics. Recent theoretical work (Elliott et al. (2022), Grossman et al. (2023)) studies optimal diversification policy under supply network disruptions. Supply chain disruptions impose significant costs, as shown by Barrot and Sauvagnat (2016) and Carvalho et al. (2021), who use natural disaster timing to estimate firm linkages in supply shock propagation. Closest to our paper, Castro-Vincenzi (2022) models how multinational automobile manufacturers reallocate production in response to climate shocks. We contribute to this literature by examining how the endogenous response to disruptions via diversification shapes competition and long-run market structure.

2 Industry Background

This section provides background information on the undersea fiber-optic cable industry. Modern undersea cables use fiber-optic technology to transmit information from one end (landing station) to another. Each cable is built with several fiber pairs (for inbound and outbound traffic), though cables typically enter service with only a fraction of capacity "lit" or activated. Unlit capacity represents spare capacity that can be brought online without laying a new cable, though incremental investment in capital equipment at either landing station is required. Lighting additional capacity is a major capital expense and takes months to complete, and once capacity is lit, it remains lit; i.e., providers do not scale capacity up and down in real-time as in electricity markets. Additional capacity can also be added by using new wavelengths for data transmission (i.e., exogenous improvements in wavelength technologies).

Laying new cables (and associated equipment) is a costly endeavor, with total construction costs in the hundreds of millions of dollars. Post-construction operational costs include electricity costs, and landing station staffing and maintenance costs in the range of $5 million or more annually, as well as cable repairs (discussed below). Electricity costs are one of the main components in a cable's variable costs. Cables are designed with a theoretical working lifespan of 20 to 30 years. Because the industry is young, only a few instances of cable exits are observed in our data and occur mostly near the end of our sample, after 2020.[2]

Undersea cable faults occur frequently, most commonly from accidental shipping interference (e.g. fishing trawlers or anchors) though interruptions from natural causes and intentional sabotage also occur. The rate of failure is significantly lower than that of terrestrial cables, though it is offset by a much higher cost of and longer length of repair: costs range from $250 thousand to $1 million, with typical time frames of one to several weeks (though repairs taking months are not uncommon). Industry estimates are that hundreds of undersea cable faults occur each year; demand for repairs to the roughly 400 active cables (in addition to the laying of new cables) is high enough to sustain a global fleet of at least 60 cableships. In order to hedge the risk posed by the possibility of cable disruptions, buyers purchase bandwidth on several cables on a given route, as indicated by various interviews we conducted with industry professionals:

> *"No customer would buy capacity on a single cable. Cables break all the time and that is very risky. Customers are buying on multiple cables."* (Alan Mauldin, TeleGeography).

Cables are owned by investors. In the majority of cases, a single investor owns the cable, but for some of the largest investments (e.g., Trans-Atlantic cables), multiple investors may form a consortium and share the investment cost. The median number of investors per cable is 1 and the mean is 2.5. On a few routes, an investor may have an ownership stake in multiple cables (e.g., AT&T on the United Kingdom–United States route), though this remains the exception rather than the rule. Investors have historically been major telecommunication

[2]Decommissioned cables include TAT-14, CANTAT-3, Tasman 2, Atlantis-2.

companies and governments; in recent years, they also include private enterprises specializing in undersea telecommunication and large technology firms. Investors appoint an operating firm (a "firm" hereafter) to administer the day-to-day operations of the cable, and we take this firm as the decision-maker in our industry model.[3] Economic regulation (of prices, provisioning of backbone traffic, interconnections, etc.) is nonexistent in this industry. Non-economic regulation (e.g., permitting requirements by the FCC to land a cable in the US or data protection policies such as GDPR in the EU) may affect entry costs and demand.

Cable firms sell bandwidth in wholesale markets, and this bandwidth is purchased by a variety of customers, which can be broken down into four broad categories: internet service providers (e.g., Comcast, Verizon), content providers (e.g., Google, Meta), private enterprises such as financial institutions or stock markets, and research institutions. Contracts are negotiated bilaterally between buyer and seller (the firm administering a cable) and specify: the duration of the lease agreement (typically a quarter to a year), the "capacity" or maximum bandwidth the lessee can utilize for data transmission across the cable, and service level agreements (SLA) specifying performance metrics such as guaranteed uptime, latency targets, and repair times in the event of a cable fault.[4]

Bandwidth is a product that is vertically differentiated by speed (also referred to as capacity), i.e., buyers can purchase leases at 10Gbps (Gigabytes per second), 40Gbps, 100Gbps, etc. Over the period studied in this paper, 10Gbps is by far the most common product purchased (100Gbps was introduced only in the last few years of our sample and 40Gbps is relatively rare), as discussed in the next section. While cables can be differentiated on other dimensions (e.g., more favorable SLA), discussions with industry professionals indicate that the scope for differentiation is limited. There exists limited within-market price dispersion reflecting different contract specifications or buyer size: e.g., large buyers such as Google pay lower prices for bandwidth than smaller buyers. Finally, some firms act both as buyer and seller: e.g., AT&T is an investor on some routes but also purchases bandwidth on other routes.[5]

---

[3] For example, the cable SEA-ME-WE 3 is owned by a consortium led by France Telecom and China Telecom, but is administered by Singtel, a Singaporean telecommunications operator.

[4] SLAs often include penalties for the operating firm if these metrics are not met and have additional details on maintenance responsibilities, processes for dealing with cable faults or damage, and any associated costs. An uptime SLA, excluding any cable fault event, might be in the range of 99.5% to 99.99% (percentage of contracted period over which the service is operational). In case of a cable fault, an SLA might specify a target of 12 to 48 hours for initial response and assessment, with a broader window (e.g., a few weeks) for complete repair.

[5] Content providers, such as Google, Meta and Amazon, have in recent years begun to invest in undersea cables either directly or by entering into a purchase agreement. In Caoui and Steck (2024), we examine this phenomenon more closely, studying 50 such investments and exploring e.g., how investment decisions are driven by the locations of content providers' data centers. However, of the 44 cables in which those 50 investments were made, 7 entered service before 2018 (before this type of investment became more

The industry shares several characteristics with the midstream oil and gas sector. Owners construct, maintain, and operate pipelines or ships (cables) and lease capacity (bandwidth) to downstream companies (e.g., internet service providers). Contracts, negotiated bilaterally, specify the pricing, quantity, duration, and terms of service. Some pipelines (cables) are owned or funded by companies primarily operating in downstream or upstream markets, while others are owned by specialists in the midstream sector.

3 Data and Descriptive Statistics

This section describes the data used in our empirical analysis: its sources, the definitions of variables of interest, and key trends in those variables.

3.1 TeleGeography

The main data used in this study are from data provider TeleGeography, a telecommunications market research firm which has been in operation since 1989. TeleGeography collects and retains comprehensive datasets on various aspects of the telecommunications industry. We make use of data from two separate datasets focusing on undersea fiber-optic telecommunications: The Global Bandwidth Research Service and the Wavelengths Pricing Suite.

The Global Bandwidth Research Service (GBRS) contains detailed data on undersea cables. The data on each cable include: landing points, date of entry into service, ownership (list of investors), designed (or potential) capacity, length, and construction costs. In addition, GBRS features a yearly panel of used bandwidth at the country-pair level. Used bandwidth describes the sum of bandwidth deployed actively in a route, it provides a measure of underlying international bandwidth demand. Finally, GBRS provide information on the location and entry date of data centers owned by major technology firms.[6]

To our knowledge, this dataset is the most comprehensive source available on the undersea cable industry. Nonetheless, some data limitations are worth noting. We do not observe cable-level revenue or market shares, nor do we have information on contracts signed between buyers and sellers. The availability of aggregate demand for each country-pair will inform the specification of our industry model, in particular, when it comes to the bandwidth demand model.

We present summary statistics from the GBRS data in Tables 1 and 2. Table 1 contains descriptive statistics on the GBRS data pertaining to undersea cable supply. The table

widespread) with 13 more entering service between 2018 and 2021; the remainder had not entered service by the end of our sample period. Given the novelty of this trend and the minimal impact it has had in our sample period, e.g., on price and bandwidth trends, we abstract from this consideration and do not model content providers' vertical integration.

[6] These are Meta, Microsoft, Amazon, and Google. The count data we use corresponds to the number of "cloud regions." A cloud region is a physical location (e.g., Palo Alto, US) where providers typically cluster multiple data centers (e.g., as of 2022, AWS operates three data centers in the Palo Alto cloud region).

contains data for the global network of undersea cables. Columns 2 and 3 present the number and total length (in km) of new cables that come online in each year; cumulative counts of the number of active cables and the total network length are shown in columns 5 and 6. Column 4 contains the (unweighted) average construction cost of the new cables in each year, normalized by the length of the cable; construction costs are reported for only 47% of cables in the data. Columns 7 through 9 contain information on potential and active capacities: potential capacity (in Tbps) is shown in column 7, and lit (or active) capacity is reported in column 8.[7] The share of capacity that is active (i.e., the utilization rate) is reported in column 9.

Several trends are visible in Table 1. The industry has experienced robust growth: the number of active cables and the total network mileage have more than doubled over this time frame. This growth is even more striking when considering the larger capacity of newer cables: both potential and lit capacity have increased by two orders of magnitude. Notably, the entry rate has remained steady despite a substantial amount of spare (unlit) capacity each year—the share of lit capacity peaks at just 35% in 2020. This suggests that capacity constraints were not a first-order concern in this industry during the sample period.

Table 1: Descriptive Statistics on Undersea Cable Supply

Year	New Cables	New km '000s	Mean cost $m / 1000km	Cum. Cables	Cum. km '000s	Potential Tbps	Lit Tbps	Share Lit
2005	13	24.9	44.3	180	572.5	87.1	10.6	0.12
2006	11	31.7	36.5	191	604.2	89.2	14.3	0.16
2007	11	7.2	80.7	202	611.4	100.5	16.5	0.16
2008	20	44.3	54.6	222	655.7	121.0	26.8	0.22
2009	17	69.1	44.0	239	724.8	142.3	34.3	0.24
2010	14	61.0	41.5	253	785.7	239.5	44.1	0.18
2011	14	29.4	39.3	267	815.1	309.8	54.5	0.18
2012	22	61.6	217.4	289	876.7	470.7	73.5	0.16
2013	11	16.9	56.6	300	893.6	748.4	93.4	0.12
2014	15	30.2	141.9	315	923.9	981.2	137.4	0.14
2015	9	19.3	61.0	324	943.1	1174.6	193.9	0.17
2016	16	64.7	466.6	340	1007.8	1479.7	292.4	0.20
2017	15	74.1	107.8	355	1081.9	1836.6	412.0	0.22
2018	18	74.0	34.7	373	1155.9	2359.3	563.8	0.24
2019	23	31.5	261.2	396	1187.4	2521.5	784.6	0.31
2020	21	68.6	103.7	417	1256.0	3171.1	1095.2	0.35
2021	17	45.3	225.3	434	1301.3	3928.3	1346.7	0.34

Note: Columns 2-4 show the number of new cables entering service in a given year, their total length (in thousands of kms), and their average length-normalized cost (in millions of USD per thousand km). Columns 5-6 show the cumulative number and total length of active cables by year. Column 7 contains the networks' potential capacity, measured in Tbps, and column 8 contains the 'lit' or activated capacity; column 9 shows lit capacity as a share of potential capacity.

[7] These figures represent *inter-regional* capacities, corresponding to the used inter-regional bandwidth total in column 5 of Table 2.

We present additional summary statistics in Table 2, shifting focus to demand-side factors. Columns 2 and 3 report the number of unique cities and countries connected to the internet via undersea cables by year. Columns 4 and 5 show international bandwidth usage totals in Tbps, with column 4 providing the global total (inter- and intra-regional) of international traffic and column 5 reporting inter-regional traffic only (comparable to the inter-regional potential and lit capacity figures in Table 1). Columns 6 and 7 provide information on the data center regions (see footnote 6) present in TeleGeography's dataset: column 6 shows the number of new data center regions opened each year (starting in 2006), while column 7 reports the cumulative total. The data in Table 2 also illustrate dramatic industry growth. The number of cities connected to undersea cables more than doubles, while the share of countries connected increases from roughly 66% to nearly 90%. The number of data center regions rises from 0 to 118, with significant growth occurring in recent years. Finally, used bandwidth grows by three orders of magnitude over the sample period.

Table 2: Descriptive Statistics on Bandwidth Demand

			Used Bandwidth (Tbps)		Datacenters	
Year	Cities	Countries	Total	Inter-Region	New	Cumulative
2005	549	126	5.0	1.8		
2006	603	131	7.0	2.6	2	2
2007	623	132	11.5	4.3	3	5
2008	665	133	19.2	6.7	1	6
2009	694	141	30.6	10.2	4	10
2010	726	148	46.8	15.9	4	14
2011	753	150	70.1	23.3	7	21
2012	820	164	101.2	34.1	6	27
2013	834	165	145.7	49.6	2	29
2014	890	167	212.5	71.2	10	39
2015	923	167	300.0	99.5	6	45
2016	949	167	443.0	145.4	12	57
2017	984	169	666.6	224.5	13	70
2018	1,040	172	997.4	343.0	12	82
2019	1,097	172	1,466.4	508.5	13	95
2020	1,152	175	2,124.8	715.5	15	110
2021	1,189	175	2,885.6	937.2	8	118

Note: Columns 1 and 2 contain the number of unique cities and countries that have an undersea cable landing point. Column 3 contains the total used bandwidth of the network by year in Tbps. Columns 4 and 5 provide counts (net new and cumulative) of datacenters contained in Telegeography's data; this data series begins in 2006 and was not available for 2021 at the time of writing.

We also acquire data on bandwidth prices from TeleGeography's Wavelengths Pricing Suite. The Wavelengths dataset contains quarterly prices collected by TeleGeography from bandwidth providers. Each observation in the dataset represents the quoted monthly price on a 1-year unprotected lease in a particular capacity segment (e.g., 10Gbps, 100Gbps) on a particular city-to-city route from a particular bandwidth provider.[8] The provider name is anonymized. Coverage is not comprehensive, but TeleGeography asserts that data is sourced on a voluntary basis from dozens of providers. In this paper, we focus our analysis on the 10Gbps capacity segment as it is by far the most commonly bought capacity over our sample period.

Table 3 presents descriptive statistics from this dataset for the 10Gbps capacity segment from 2005 to 2021. Column 2 presents the number of unique price quotes each year, and columns 3-5 contain the 25th, 50th and 75th percentile of quoted prices (monthly rate, in thousands of US$). Table 3 reveals a few empirical features worth noting. First, prices fall significantly over time at all reported percentiles. Second, significant price dispersion exists within a year, with the ratio of 75th to 25th percentile prices ranging from roughly 2 to more than 10. This price dispersion is expected given the heterogeneity in costs and distances across markets (e.g., U.S.-Japan versus U.K.-France). Third, the rise in the number of cities connected by cables (shown in Table 2) is reflected in the increase in the number of price quotes available.

3.2 Supplementary Data

In addition to the data from TeleGeography, we use several auxiliary data sources in our estimation approach. A first set of auxiliary data is sourced from the CEPII Gravity Database, which provides macroeconomic and trade-related variables suitable for estimating the determinants of international trade (Conte et al. (2021)). We source data on GDP, trade flows, distances, and whether two countries share a common language or border.

Another data series used in estimating demand is broadband subscriptions, which we use as a proxy for internet penetration. This data is accessed from the World Bank, which maintains an annual panel tracking fixed broadband subscriptions with speeds of at least 256Kb per second. This includes residential and business customers but excludes mobile internet users.

When estimating demand for bandwidth, we use an instrumental variable strategy with electricity prices as cost-shifters. Panels of electricity generation shares are sourced from Our World in Data (ourworldindata.org), with detailed country-level data on electricity generated by coal, gas, and oil. Price series for coal, gas, and oil are sourced from the Federal Reserve Economic Data (FRED) (https://fred.stlouisfed.org/). We use the Brent crude series for oil, the global Australian coal price, and the global EU price of natural gas; the latter two are converted to dollars per barrel of oil equivalent (BOE).

[8] Protected bandwidth refers to leased capacity on a cable backed by a secondary "protection" cable. If the primary cable fails, traffic is automatically rerouted, minimizing downtime. Unprotected bandwidth lacks this backup, making it more susceptible to outages. Most traded bandwidth on undersea cables is unprotected.

Table 3: Descriptive Statistics on Bandwidth Prices (in thousands of US$)

		10 Gbps Price Percentiles		
Year	Quotes	25th percentile	50th percentile	75th percentile
2005	359	31.8	55.5	75.9
2006	867	13.5	22.0	38.3
2007	1,437	12.0	15.0	23.3
2008	2,510	10.0	13.0	18.0
2009	1,681	8.9	11.9	17.1
2010	1,830	7.2	10.9	19.8
2011	2,334	5.5	9.4	45.0
2012	2,431	5.0	9.0	50.0
2013	2,531	4.1	6.6	42.1
2014	2,555	3.3	6.0	45.0
2015	2,512	2.6	5.0	30.0
2016	2,967	2.1	4.5	28.0
2017	3,547	1.1	3.2	24.0
2018	3,623	1.1	3.1	19.9
2019	3,602	1.2	2.1	16.2
2020	3,124	1.3	2.4	15.0
2021	3,309	1.1	2.2	12.0

Note: Columns 2-5 show data for the 10Gbps service. Each row contains the number of unique posted prices, as well as the 25th, 50th, and 75th percentiles of posted prices. Prices represent monthly prices for a 1-year unprotected lease.

Finally, we compile data on cable faults. TeleGeography's GBRS dataset reports publicly disclosed major faults since 2016 with detailed information (e.g., date fault discovered, start of repairs, length of repairs, cause). We supplement this by hand-collecting fault data for 2013-2016 from newspapers, resulting in a final sample of 168 faults. Table A8 presents the yearly fault counts.

The dataset represents only a subset of total faults, as noted by industry experts. This is because the data collection focuses on publicly disclosed or notable faults. Nevertheless, it provides valuable insights into disruption incidents and repair durations. We use the cable faults data in our estimation: (1) to estimate the distribution of disruption shocks and downtime, and (2) to verify that our estimated cable operating costs align with fault propensity. We also use the fault data in robustness checks as an exogenous shifter of entry costs and market structure in demand estimation.

These data show that disruptions are often geographically isolated, enabling diversification. While rare cases involve more than two cables disrupted in a region-quarter,[9] over 78% of faults involve only a single cable (Appendix Figure A6). Most country pairs are connected by multiple cables, providing scope for diversification (Appendix Figure A7).[10]

[9] Table A2 presents the list of region pairs.

[10] Two country pairs have zero cables in 2021 due to cable exits in 2020.

3.3 Market Definition

In anticipation of the industry model presented in the next section, we discuss our market definition. As outlined above, our primary data for estimating demand comes from TeleGeography's GBRS and Wavelength Pricing Suite.

The GBRS dataset contains data on used bandwidth at a country-pair level on an annual basis. The Wavelength Pricing dataset contains price quotes on bandwidth at a city-to-city level on a quarterly basis. To strike a balance between the granularity of the pricing data relative to the used bandwidth data, we combine the two datasets as follows. First, we aggregate price quotes from the city-to-city level to the country-to-country level (we use a weighted-average specification where the weights are the number of city-to-city price quotes, but have also experimented using simple average and median); we also multiply the quoted monthly lease prices by three to convert them from a monthly to a quarterly price. Second, we interpolate the annual bandwidth used to the quarterly level.

At the end of this process, a market-period in our estimation sample is a country-to-country pair in a calendar quarter. Unsurprisingly this panel is not perfectly balanced: cables enter and price quotes are collected for new country pairs over time. Details on the sample of markets are presented in Appendix A.[11]

4 Industry Model

Time is discrete with an infinite horizon $t = 0, 1, 2, ...$, and a period corresponds to a calendar quarter. Markets are assumed to be independent of each other. In what follows, we consider a specific market m.

4.1 Bandwidth Demand Under Supply Disruptions

We start by presenting static demand-side objects that are not determined by the dynamic equilibrium. The objective of this section is to derive the aggregate demand for bandwidth faced by cable operators from a consumer-level utility maximization problem under supply disruptions. The m and t subscripts are omitted in this section. Proofs are included in Appendix B.

A representative buyer faces n symmetric firms (i.e., cable operators) supplying bandwidth, a homogeneous product. A firm operates one cable per market. Cables are subject to disruptions: if the buyer purchases bandwidth q_i from firm i, she receives bandwidth $\delta_i q_i$, where δ_i is a random variable with support $[0, 1]$, mean μ, and a density $g(\cdot)$ that is continuous and strictly positive over the support. Denote the corresponding cumulative distribution

[11] Some markets consist of contiguous countries (e.g., France-Spain, US-Canada). To address concerns that terrestrial cables may carry significant bandwidth in these markets, we either control for contiguity in demand estimation or exclude the 34 markets with contiguous countries from our analysis.

function $G(\cdot)$. For simplicity, we assume that disruption events are *iid* across firms.

We break down the buyer's optimization problem into two steps. In the first step, the buyer maximizes her utility by choosing a total amount of *used* bandwidth B. The bandwidth B corresponds to bandwidth the buyer is committed to using—either for its internal operation if the buyer is a content provider or selling to its downstream customers if the buyer is an ISP. In the second step, the buyer provisions a level of *purchased* bandwidth Q, in order to minimize her expected cost of disruptions; these costs are incurred whenever the realized purchased bandwidth falls short of the used bandwidth B. We expect $Q > B$, i.e., the buyer builds a buffer stock in order to insure against the cost of disruptions. We solve the buyer's problem by backward induction.

Bandwidth Provisioning Problem. To ease notation, we assume that the buyer splits the purchased bandwidth equally across firms, that is, for all i, $q_i = q$ and $Q = nq$.[12] Given used bandwidth B, the buyer chooses the purchased bandwidth Q to solve the following cost minimization problem

$$\min_{Q} \quad PQ + \gamma P\,\mathbb{E}\left[\max\left\{B - Q\frac{1}{n}\sum_{i=1}^{n}\delta_i, 0\right\}\right] \quad \text{s.t} \quad Q \geq B\,, \tag{1}$$

where P is the price of bandwidth. The buyer chooses purchased bandwidth Q to minimize the sum of the direct cost of purchasing bandwidth (first term) and the disruption cost (second term). The buyer incurs a disruption cost whenever the realized bandwidth $Q\frac{1}{n}\sum_{i=1}^{n}\delta_i$ falls short of the used bandwidth B. The cost of disruptions (per Gbps and calendar quarter) γP is (without loss of generality) proportional to the price of bandwidth and γ is assumed to be greater than one.

This specification of the disruption cost captures several scenarios: e.g., it can represent the reputational cost to an ISP of failing to provide the bandwidth B they have committed to supplying to their downstream customers, the cost of rerouting traffic via alternative paths at the last minute, the financial costs from delayed transactions for stock exchanges, or the lost profits to a content provider providing cloud computing services to their downstream customers. Importantly, the disruption cost does not include the repair cost of the cable, which is incurred by the firm. In the case of undersea cables, repair costs are negligible compared to the disruption costs to buyers (APEC (2012)).

To guarantee an interior solution, we assume that the cost of disruption is high enough,

[12] This assumption is justified given that, in this industry, transaction costs from negotiations with multiple suppliers are arguably negligible compared to the size of contracts.

that is:[13]

$$\frac{1}{\gamma} < \mu\,. \tag{2}$$

As shown below, this condition ensures that $Q > B$, so that the constraint in Problem (1) is not binding. Let $g_n(\cdot)$ denote the density of the random variable $\frac{1}{n}\sum_{i=1}^{n}\delta_i$ (with support over $[0,1]$), and $G_n(.)$ denote the corresponding cumulative distribution function. The first-order condition with respect to Q is:

$$\frac{1}{\gamma} = \int_0^{\frac{B}{Q}} u g_n(u) du\,. \tag{3}$$

Lemma 1. *The cost minimization problem has a unique solution, with $Q > B$. Denote the cost-minimizing ratio of used-to-purchased bandwidth as*

$$\left(\frac{B}{Q}\right)^* \equiv \widetilde{f}(n,\gamma)$$

The ratio $\widetilde{f}(n,\gamma)$ is strictly decreasing in γ.

Lemma 1 shows that the cost-minimizing ratio of used-to-purchased bandwidth exists and derives comparative statics with respect to γ. When disruptions are more costly, i.e., γ increases, the buyer purchases more bandwidth Q to better insure against disruption risk.

At the optimum, the cost function (1) can be written as a function of used bandwidth B

$$B \cdot P \cdot \left(\frac{1}{\widetilde{f}(n,\gamma)} + \gamma\,\mathbb{E}\left[\max\left\{1 - \frac{1}{\widetilde{f}(n,\gamma)}\frac{1}{n}\sum_{i=1}^{n}\delta_i, 0\right\}\right]\right) \equiv B \cdot P \cdot \widetilde{h}(n,\gamma). \tag{4}$$

We impose the following restriction on the distribution of disruption shocks.

Assumption 1. ***(Single-crossing)*** *For all $n \geq 1$, there exists a unique $x_n \in (0,1)$, such that,* (1) $G_n(x_n) = G_{n+1}(x_n)$*,* (2) $G_n(x) > G_{n+1}(x)$ *if $x < x_n$, and* (3) $G_n(x) < G_{n+1}(x)$ *if $x > x_n$.*

Assumption 1 is satisfied by commonly used distributions: e.g., symmetric distributions (e.g., normal, uniform), and more generally, unimodal distributions or distributions with monotone density (e.g., Exponential, Gamma, Beta with parameters outside of the unit square).[14] By the Central Limit theorem, any distribution will satisfy Assumption 1 for n large enough. However, we are most interested in deriving results for small n. Assumption 1 allows us to discipline the behavior of G_n for small n, ruling out non-standard distributions

[13] This assumption is satisfied in practice because μ is very close to one, whereas γ is at least an order of magnitude greater than one. See the model calibration in Appendix C.

[14] In general, for distributions with support other than $[0,1]$, the assumption holds on the interior of the support.

(e.g., multi-modal or with sharp peaks in their density function). In practice, this assumption is satisfied by the distribution of disruption shocks in our particular application.

Proposition 1. *Under Assumption 1, the sequence $\widetilde{h}(n,\gamma)$ is strictly decreasing in n. Moreover, this sequence converges to $\frac{1}{\mu}$.*

Proposition 1 shows that, as the number of firms increases, the expected total cost of provisioning bandwidth B (Equation (4)) decreases, holding the bandwidth price P fixed.[15] Intuitively, with more suppliers, the variance in realized bandwidth decreases, and the buyer is better able to diversify risk; she can reduce over-provisioning. As n grows to infinity, the realized purchased bandwidth approaches μQ, the purchased bandwidth approaches B/μ, the expected disruption cost converges to zero, and therefore, the total cost converges to $PQ = PB/\mu$.

Bandwidth Usage Problem. Given the solution to the cost minimization problem, the buyer chooses used bandwidth B to maximize a concave utility (net of the bandwdith cost)

$$\max_{B} \frac{1}{\alpha+1} B^{\alpha+1} - BP\widetilde{h}(n,\gamma)\,, \tag{5}$$

for $-1 < \alpha < 0$. The first-order condition with respect to B is:

$$B = P^{\frac{1}{\alpha}} \widetilde{h}(n,\gamma)^{\frac{1}{\alpha}}\,. \tag{6}$$

The logarithm of used bandwidth can be expressed as a function of the bandwidth price and the number of cables as follows:

$$\log(B) = \frac{1}{\alpha}\log(P) + \frac{1}{\alpha}\log\left(\widetilde{h}(n,\gamma)\right). \tag{7}$$

An important corollary of Proposition 1 is that the function $\frac{1}{\alpha}\log\left(\widetilde{h}(n,\gamma)\right)$ is increasing in n (recall $\alpha < 0$). That is, with supply disruptions, aggregate demand is increasing in the number of suppliers n: by allowing the buyer to better diversify disruption risk, entry of additional suppliers expands demand outward. This aspect of the model shares similarities with models where consumers value product variety (e.g., Spence (1976), Mankiw and Whinston (1986)). Relative to previous empirical work on this topic, which typically relies on discrete choice models (e.g., nested logit, random coefficient logit), the welfare effect of new entry in our model are derived from a consumer-level risk diversification problem, and are not driven by ex-ante restrictions on the unobservable characteristics space (Ackerberg and Rysman (2005)).

[15] The term $\widetilde{h}(n,\gamma)$ can be interpreted as an "insurance premium," that is, the extra cost beyond the direct cost $B \cdot P$ that the buyer incurs in order to minimize her expected disruption costs. Proposition 1 states that this insurance premium is decreasing in n.

When taking this model to the data, we assume that used bandwidth in market m in period t (measured in Gbps), denoted B_{mt}, is a function of the bandwidth price P_{mt} (in \$US), the number of cables operating in the market n_{mt}, and an exogenous demand state denoted d_{mt}, which captures various market and time specific demand factors such as the level of internet penetration or the number of data centers operating in the market.[16] The aggregate demand curve takes the log-linear form:

$$B_{mt}(d_{mt}, P_{mt}, n_{mt}) = \exp\left(d_{mt} + \sum_{n} \alpha_n \mathbb{1}\{n_{mt} = n\}\right) P_{mt}^{\alpha_p}. \tag{8}$$

Three points are worth noting. First, the bandwidth provisioning problem provides a direct mapping between the reduced-form aggregate demand in Equation (8) and structural parameters in Equation (6), which are rooted in consumer preferences, the distribution of disruption shocks, and disruption costs. In particular, the price elasticity α_p is related to the curvature parameter α, whereas the coefficients α_n capturing the effect of n on $\log(B)$ are a function of the curvature parameter α, the distribution of δ_i, and the disruption cost γ. Second, if the buyer cannot diversify risk by sourcing Q from multiple firms (i.e., setting n equal to one in Problem (1)), one can obtain the counterfactual aggregate demand curve by replacing $\widetilde{h}(n, \gamma)$ with $\widetilde{h}(1, \gamma)$ in Equation (7), or alternatively by setting n_{mt} to one in Equation (8). Third, if disruptions are shut down (i.e., $\delta_i = 1$, almost surely), then $\widetilde{f}(n, \gamma)$, and therefore, $\widetilde{h}(n, \gamma)$ equal one; $\log(B)$ no longer depends on n.

In Appendix C, we compute the function $\frac{1}{\alpha} \log\left(\widetilde{h}(n, \gamma)\right)$ given the empirical distribution of cable disruptions, and shows that B_{mt} and Q_{mt} are both increasing in n_{mt}, with decreasing marginal effects. The appendix also provides evidence that predictions of the bandwidth provisioning problem match used-to-purchased bandwidth ratios observed in the data.

While the buyer's utility is a function of used bandwidth B_{mt}, we note that transfers from buyers to firms and firm profits depend on purchased bandwidth Q_{mt} (which equals $B_{mt}/\widetilde{f}(n, \gamma)$). Firms' variable costs are a function of purchased bandwidth.[17] In what follows, we take care to distinguish between B_{mt} and Q_{mt} when calculating consumer surplus and firm profits.

4.2 Cable Entry Decisions

Investment in undersea cables and bandwidth supply are modeled as a dynamic game played in each market. A player is an operating firm administering a cable. Undersea cables are durable equipment, making the decision to build a cable a typical case of investment under

[16] Market-level heterogeneity (d_{mt}) can be accommodated in the bandwidth usage problem by scaling the utility function, e.g., $\frac{1}{\alpha+1} B^{\alpha+1} e^{-\alpha d_{mt}}$.

[17] The primary variable costs like electricity consumption are determined by the cable's operational capacity and power requirements, which are constant regardless of the actual data transmission.

uncertainty. The central part of the model specifies how firms make their entry decisions as a function of market structure (i.e., the number of incumbent firms), and market-level demand and cost state variables. Finally, an equilibrium of the dynamic game is specified. In this setting, the demand and cost state transition processes are nonstationary; hence, we characterize a nonstationary Markov Perfect equilibrium of the game, where strategies and transition functions are indexed by time. We abstract from modelling the formation of investor consortia and common ownership issues.[18]

Players and States. A finite number N of symmetric firms are indexed by $i \in \{1, \ldots, N\}$. In any period t, each firm is either active in the market or a potential entrant. A firm is defined as active if it operates an undersea cable. We assume that a firm can operate at most one cable per market.[19] Therefore, we refer to a firm and a cable interchangeably. We denote its state by $s_{it} \in \{0, 1\}$, where s_{it} equals 1 if the firm is active and zero otherwise. The industry state is the aggregation of firm states $s_t \equiv \{s_{it}\}_{i \in N}$ and we let $s_{-it} = \{s_{jt}\}_{j \neq i}$ denote the state of firms other than i, and $n_{mt} = \sum_{i=1}^{N} s_{it}$ the total number of firms operating in the market.

The exogenous characteristics of the market in period t consist of two main components: (1) an aggregate demand state, denoted d_{mt}, introduced above in Section 4.1; (2) an aggregate cost state h_{mt}, which captures supply factors such as input costs (e.g., electricity), improvements in bandwidth capacity provisioning (through technological advances), the physical and geographic characteristics of the market that determine the required cable length and number of repeaters, or the frequency of cable faults. Demand factors change over time due to population and economic growth and increasing internet access and digitization. Electricity costs evolve over time due to changes in global energy prices (e.g., of oil, gas, and coal) and country-specific changes in their electricity-generation fuel mix.

[18] A cable is built if the discounted sum of expected payoffs is higher than realized entry costs and we do not model how these costs are shared between investors. Common ownership, whereby an investor has an ownership stake in multiple cables in *the same market*, is relatively rare (investors with multiple ownership stakes typically have those stakes in cables in different markets). Including a consortium formation game and common ownership considerations in the model would substantially increase its complexity and the computational burden in estimation, while not adding much explanatory power in understanding the role of diversification by buyers. To wit, the mean number of ownership stakes per cable-market is 1.1 and the 75th percentile is 1.

[19] We rule out cases where a single firm operates multiple cables in a given market for computational simplicity and because this is rare (see Section 2). To the extent we overstate competition in a few cases, we would slightly overestimate firms' marginal costs and underestimate the scale of firm-specific shocks θ^{ϵ}. Predicting a directional change in our counterfactual results from relaxing this assumption is challenging because: *i*) the multi-cable firm's profit maximization problem internalizes both business-stealing (negative) and diversification (positive) externalities. The second counterfactual exercise solves a social planner problem and relies less on the market conduct assumption. A fuller characterization of equilibrium in such markets with multi-cable owners is an interesting avenue for future research.

The vector of public information variables includes the industry state and exogenous market characteristics. All these variables are publicly observed and denoted by the vector $\mathcal{M}_t = (n_{mt}, d_{mt}, h_{mt})$.

Actions. Every period t, firms decide simultaneously and independently whether to be active or not in the market. Let $a_{it} \in \{0, 1\}$ be the binary indicator of the firm i's decision at period t, such that $a_{it} = 1$ if an incumbent firm decides to remain active in the market at the end of period t or a potential entrant decides to enter, and $a_{it} = 0$ if an incumbent exits or a potential entrant stays out at the end of period t. Firms that exit or potential entrants that decide to stay out are replaced by a new set of potential entrants in the following period.[20] We use the variable a_{-it} to denote the vector of actions taken in period t by all firms except firm i.

Firms' choices are dynamic because of partial irreversibility in the decision to enter a market, i.e., sunk costs. At the end of period t, firms simultaneously choose their action a_{it} which determines their next-period state s_{it+1}. We model the choice of entry and exit as a game of incomplete information, so that each firm i has to form beliefs about other firms' entry and exit choices.

4.3 Period Profits

Firm i's period profits, net of private information shocks, are

$$\pi_i(a_{it}, \mathcal{M}_t) = VP_i(\mathcal{M}_t) - FC_{mt}(s_{it}) - EC_{mt}(s_{it}, a_{it}) + EV_m(s_{it}, a_{it}), \tag{9}$$

where $VP_i(\mathcal{M}_t)$ are variable profits, FC_{mt} is the fixed cost incurred by firm i to operate an undersea cable in market m, EC_{mt} is the entry or set-up cost of a new cable, and EV_m is the scrap value from retiring an exiting cable. The variable profits $VP_i(\mathcal{M}_t)$ are equal to the difference between revenue from selling bandwidth on the undersea cable and the variable costs of operating the cable (e.g., maintenance, energy costs). Bandwidth is a high-tech commodity with limited scope for differentiation, and hence we assume Cournot competition among undersea cable carriers that are active in period t. We focus on a Nash equilibrium in the spot market for bandwidth. Thus, the market structure (summarized by the industry state n_{mt}), the aggregate states (d_{mt}, h_{mt}), and the aggregate demand model for purchased bandwidth (Q_{mt}) completely determines each firm's equilibrium variable profit $VP_i(\mathcal{M}_t)$ from competing in period t. This parsimonious formulation allows us to handle

[20] This assumption is for computational convenience, as we avoid having to solve an optimal stopping problem for potential entrants. It does not affect the (static) estimates of demand and cost states, although it is likely to have an effect on estimates of some of the dynamic parameters. In this industry, potential entrants face time-limited opportunities due to factors like financing windows, landing rights, permits, and environmental assessments. These factors, with their inherent time constraints, support the assumption of short-lived potential entrants.

the dynamic oligopoly game of entry and exit with a tractable state space. The fixed cost of an active firm, FC_{mt}, is incurred in any period where the firm is active ($s_{it} = 1$) and reflects the need for continual investment in maintenance, expected cable repairs, and upgrades to facilities and equipment.

At the beginning of period t, each firm draws a vector of private information shocks associated with each possible action $\boldsymbol{\epsilon}_{it} = \{\epsilon_{it}(a)\}_{a\in\{0,1\}}$. We assume that the shocks ϵ_{it} are independently distributed across firms and over time. In our application, these shocks will be distributed Type 1 extreme value, scaled by a parameter θ^{ϵ}_{m} which can vary by market.

It will be convenient to distinguish two additive components in the period profit function:

$$\Pi_{it}(a_{it}, \mathcal{M}_t, \boldsymbol{\epsilon}_{it}) = \pi_i(a_{it}, \mathcal{M}_t) + \epsilon_{it}(a_{it})\,. \tag{10}$$

4.4 Dynamic Optimization and Equilibrium

Firms make their dynamic discrete choices of entry and exit to maximize their discounted sum of expected profits. They discount their future stream of profits by a factor $\beta \in (0, 1)$, with rational expectations regarding the endogenous evolution of market structure and the exogenous evolution of demand and production costs.

We focus on Markov-Perfect Bayesian Nash Equilibria (MPE). We first define firm strategies, value functions, and then the equilibrium conditions. A firm's strategy, at time t, depends only on its payoff-relevant state variables $(\mathcal{M}_t, \boldsymbol{\epsilon}_{it})$. A strategy profile is denoted

$$\alpha = \{\alpha_{it}(\mathcal{M}_t, \boldsymbol{\epsilon}_{it}))\}_{i\in I, t\geq 0}\,.$$

Given strategy profile α, firm i's value function satisfies

$$V^{\alpha}_{i,t}(\mathcal{M}_t, \boldsymbol{\epsilon}_{it}) = \max_{a_{it}\in\{0,1\}} \left\{ v^{\alpha}_{i,t}(a_{it}, \mathcal{M}_t) + \epsilon_{it}(a_{it}) \right\}\,, \tag{11}$$

where $v^{\alpha}_{i,t}(a_{it}, \mathcal{M}_t)$ are choice-specific value functions. The choice-specific value function for active firms are given by

$$v^{\alpha}_{i,t}(a_{it}, \mathcal{M}_t) = \begin{cases} \pi_i(1, \mathcal{M}_t) + \beta\, \mathbb{E}_t \left[V^{\alpha}_{i,t+1}\left(\mathcal{M}_{t+1}, \boldsymbol{\epsilon}_{i,t+1}\right)\right] & \text{if} \quad a_{it} = 1 \\ \pi_i(0, \mathcal{M}_t) & \text{if} \quad a_{it} = 0 \end{cases}\,, \tag{12}$$

where the next-period state $\mathcal{M}_{t+1}$ is formed of the next-period market structure $n_{m,t+1}$, and exogenous market-level variables $(d_{m,t+1}, h_{m,t+1})$; the fixed costs and scrap value are included in $\pi_i(a_{it}, \mathcal{M}_t)$ as specified in Equation (9). The distribution over next-period states $\mathcal{M}_{t+1} = (n_{m,t+1}, d_{m,t+1}, h_{m,t+1})$, conditional on current state $\mathcal{M}_t$ and action profile $a_t = (a_{it}, a_{-it})$, is given by the transition matrix $F_t(\cdot|\cdot)$ which is indexed by time because the processes are nonstationary. For potential entrants, the choice-specific value functions are given by

$$v^{\alpha}_{i,t}(a_{it},\mathcal{M}_t)=\begin{cases}-EC_{mt}+\beta\,\mathbb{E}_t\left[V^{\alpha}_{i,t+1}\left(\mathcal{M}_{t+1},\boldsymbol{\epsilon}_{i,t+1}\right)\right] & \text{if} \quad a_{it}=1\\ 0 & \text{if} \quad a_{it}=0\end{cases}. \tag{13}$$

A MPE is characterized by a strategy profile α^* such that for every player, state, and period

$$\alpha^*_{i,t}(\mathcal{M}_t,\epsilon_{it})=\arg\max_{a_{it}\in\{0,1\}}\left\{v^{\alpha^*}_{i,t}(a_{it},\mathcal{M}_t)+\epsilon_{it}(a_{it})\right\}, \tag{14}$$

and beliefs about rivals' entry and exit decisions are dictated by α^*. The probability that firm i chooses action a_{it} in period t given state $\mathcal{M}_t$ (hereafter, the conditional choice probability or CCP) is defined as

$$P^{\alpha}_t(a_{it}|\mathcal{M}_t)\equiv\Pr(\alpha^*_{i,t}(\mathcal{M}_t,\epsilon_{it})=a_{it}|\mathcal{M}_t)\,. \tag{15}$$

One can express the choice-specific value function as a function of CCPs instead of strategies. That is,

$$v^{\mathbf{P}}_{i,t}(a_{it},\mathcal{M}_t)=\pi_i(a_{it},\mathcal{M}_t)+\beta\sum_{a_{-it}}\sum_{\mathcal{M}_{t+1}}\overline{V}^{\mathbf{P}}_{i,t+1}\left(\mathcal{M}_{t+1}\right)F_t(\mathcal{M}_{t+1}|\mathcal{M}_t,a_t)P_t(a_{-it}|\mathcal{M}_t)\,, \tag{16}$$

where $a_t=(a_{it},a_{-it})$ and $\overline{V}^{\mathbf{P}}_{i,t}$ is the *ex-ante* value function expressed before the realization of the private shock $\boldsymbol{\epsilon}_{it}$

$$\begin{aligned}\overline{V}^{\mathbf{P}}_{i,t}(\mathcal{M}_t)=\int\max_{a_{it}\in\{0,1\}}\Bigg\{&\pi_i(a_{it},\mathcal{M}_t)+\epsilon_{it}(a_{it})\\ &+\beta\sum_{a_{-it}}\sum_{\mathcal{M}_{t+1}}\overline{V}^{\mathbf{P}}_{i,t+1}\left(\mathcal{M}_{t+1}\right)F_t(\mathcal{M}_{t+1}|\mathcal{M}_t,a_t)P_t(a_{-it}|\mathcal{M}_t)\Bigg\}\,dG(\epsilon_{it})\,.\end{aligned} \tag{17}$$

If private shocks are distributed Type 1 extreme value (with scale parameter θ^{ϵ}_m), an optimal strategy for firm i will map into equilibrium CCPs of the form

$$P_t(a_{it}|\mathcal{M}_t,\mathbf{P})=\frac{\exp\left(\frac{v^{\mathbf{P}}_{i,t}(a_{it},\mathcal{M}_t)}{\theta^{\epsilon}_m}\right)}{\sum\limits_{a'\in\{0,1\}}\exp\left(\frac{v^{\mathbf{P}}_{i,t}(a',\mathcal{M}_t)}{\theta^{\epsilon}_m}\right)}\,. \tag{18}$$

Multiple equilibria may exist. In our identification approach, we assume that markets with the same observable characteristics select the same equilibrium, and we verify that estimates obtained using a two-step estimator—a procedure that is robust to multiplicity—are close to those obtained using iterative methods. For our first counterfactual exercise, we initialize the algorithm at a large number of starting values and converge systematically to a unique fixed point. For the second counterfactual exercise, we solve a social planner problem (as a single-agent dynamic decision problem).

5 Identification and Estimation Approach

5.1 Identification

We follow the literature on the identification of dynamic decision problems and assume that the discount factor and the distribution of firm shocks are known.[21] We allow the scale of the firm shocks θ_m^ϵ to depend on entry costs, as described below. The scale of firm-specific shocks is identified because variable profits are treated as observed when estimating the dynamic model.

Aguirregabiria and Suzuki (2014) study the identification of market entry and exit games. They show that fixed costs, entry costs, and exit values are not separately identified. In our model, we normalize the exit value to zero and estimate entry and fixed costs. Industry reports suggest that while raw materials and equipment from retired cables can be reused, most (94%) unused undersea *telephone* cables are abandoned on the seabed The Guardian (2016), supporting our normalization of scrap values.

The industry's recent nature presents us with a challenge: exit events are rare. Without observing exits, we cannot separately identify entry costs from fixed costs, even with the above normalization. To address this, we use cable-level construction cost data to estimate entry costs outside the dynamic model. Once entry costs are estimated, firms' optimal entry decisions allow us to recover fixed costs.

Two key parameters enter our estimation and counterfactual analysis: the elasticity of bandwidth demand with respect to prices and the dependence of bandwidth demand on the number of operating cables. The first is identified using an instrumental variable approach based on cost shifters, detailed in Section 5.3. The second set of parameters is identified under assumptions similar to the endogenous product variety literature (e.g., Eizenberg (2014), Wollmann (2018), Fan and Yang (2020)): firms make entry decisions before current-period transient demand and cost shocks are realized. This assumption is natural in our setting, as entry decisions typically occur at least a year before the cable becomes operational, and timing of entry—after conditioning on state variables—depends on idiosyncratic shocks (e.g., construction delays) unrelated to transient demand shocks. Robustness checks for this assumption are provided in Section 5.3, and we discuss the role of time-to-build in Appendix F.1.

Finally, given the availability of a long panel, we control for persistent market-level unobserved heterogeneity in demand and costs with fixed effects.

[21] In the estimation approach, we experiment with discount factors ranging from 0.90 to 0.975, with 0.95 as a baseline, and assume that firm-specific shocks follow a Type-1 extreme value distribution.

5.2 Solution Method and Estimation Approach

In our empirical approach, we follow three steps. First, we estimate demand for bandwidth in each market to recover price elasticities and the relationship between demand and the number of operational cables. Second, we recover the marginal costs of bandwidth implied by the demand estimates and the firms' first-order profit-maximization conditions. These static estimates for each market-period enable computation of variable profits per cable under different market structures and time periods. We also estimate the transition processes of exogenous components affecting demand and cost states. Third, we incorporate these profits into a dynamic game of entry and competition among cable operators to estimate the fixed costs of operating cables and the scale of firm-specific logit shocks.

In this section, we take the outputs of the first two steps as given and give an overview of the methods used to solve and estimate the dynamic game (third step). The methods used are based on existing literature, with full details presented in Appendix D. We begin by discussing how the model is solved, as estimation of the dynamic model involves only small extensions to the solution procedure.

Solution Method. The dynamic game is solved via policy iteration (Judd (1998), Rust (2000)). This approach consists in iterating repeatedly between two steps: a given iteration starts by updating the ex-ante and choice-specific value functions given the current vector of CCPs (policy evaluation), then these value functions are used to update the vector of CCPs (policy improvement). The algorithm iterates until value functions and CCPs converge, up to a pre-defined tolerance level.

Three important features are worth highlighting. First, because the demand and costs states may follow non-stationary processes, we adopt as equilibrium concept a symmetric non-stationary MPE, in which time becomes a state variable. To maintain tractability, we assume that the industry enters a stationary regime after some period T (in practice, we use the last quarter of 2021). Second, because there are too few exit events over our sample period, we do not model exit decisions. Cable operators choose optimally when to enter, but once in the market, we assume that incumbents remain active. We revisit this assumption and incorporate cable design life (in the order of 25 years) and exogenous retirement dates in Appendix F.2. Third, we assume that, among all potential entrants, only one firm has an opportunity to enter each period.[22] This assumption serves two purposes: it rules out the possibility of multiple cables entering in the same period which does not occur in our data; and as detailed in Section 6, it greatly simplifies the solution of the social planner's problem.

Estimation Approach. The objective of the dynamic game estimation is to recover the

[22] The total number of firms N per market is set to the maximum number of cables observed in the data for that market plus two.

level of fixed costs FC_m and the scale parameter of the firm-specific shock θ^{ϵ}_m. We denote these parameters $\boldsymbol{\theta}_m \equiv (FC_m, \theta^{\epsilon}_m)$. Our baseline estimator is a nested pseudo-likelihood (NPL) procedure following Aguirregabiria and Mira (2007).[23] This procedure relies on a similar iterative procedure as the one used to solve the model, with an added estimation step.

In some cases, the NPL algorithm may fail to converge if the fixed point corresponding to the data generating process is unstable (Pesendorfer and Schmidt-Dengler (2010)). To address this concern, we implement several alternative estimators: two-step estimators (1-PML and 1-MD) and the spectral algorithm recently proposed by Aguirregabiria and Marcoux (2021). The latter estimator does not iterate on the best-response mapping to attain a fixed-point, rather, it solves for the root of a nonlinear system of equations using a quasi-Newton method. As a consequence, the spectral approach is able to find unstable fixed points that would be unattainable by the NPL algorithm. One advantage of the iterated procedures above (e.g., imposing equilibrium restrictions) is that they are robust with respect to the consistency of the initial estimates of CCPs.

We compute standard errors by bootstrap where markets are re-sampled (300 replications). To account for the uncertainty in the static demand and marginal cost estimates, the entire three-step procedure (demand, marginal costs, and dynamic parameter estimation) is performed on each bootstrap sample.

5.3 Demand Model

In this section, we estimate demand-side objects that are not determined by the dynamic equilibrium: the demand for bandwidth and the transition of the exogenous demand state variables.

Bandwidth Demand. The empirical analog of the aggregate demand curve presented in Section 4.1 (Equation (8)) is

$$\log(B_{mt}) = \alpha_p \log(P_{mt}) + \sum_n \alpha_n \mathbb{1}\{n_{mt} = n\} + \alpha_2 \mathbf{X}_{mt} + \eta_m + \eta_{r(m)t} + \epsilon_{mt}\,, \tag{19}$$

where $\mathbf{X}_{mt}$ are demand shifters, η_m is a market-specific fixed effect, and $\eta_{r(m)t}$ is a region-pair by year fixed effect. To capture potential non-linear effects of the number of cables n_{mt}, we include it as a categorical variable.[24] The term $\mathbf{X}_{mt}$ includes fixed broadband subscriptions, measures of Gross Domestic Product (GDP), aggregate trade flows, and the number of data centers; we restrict the coefficients on these variables to be the same for

[23] A "one-step ahead" CCP estimator like that proposed in Arcidiacono and Miller (2011) would help in estimating the game with a nonstationary environment but is precluded in practice by the very limited number of exit events in the data.

[24] The omitted category is "zero undersea cables." Including observations for market where undersea cables have not yet entered allows us to control for the baseline level of demand through other indirect paths.

both countries forming the market. It also includes bilateral measures such as the distance between the two countries, and whether they share a common language and/or land border (in specifications without market fixed effects). Market fixed effects capture unobserved heterogeneity such as cable fault propensity. Region-pair by year fixed effects allow us to capture regional time trends, e.g. dynamically changing composition of buyers (internet backbone providers, content providers, private enterprises), or transient regional demand shocks: examples include GDPR regulation in the EU that may affect data flows between the US and Europe or the growth in content delivery networks (CDNs) that allow local storage of data near end-users.[25]

We address the issue of endogeneity of prices via instrumental variables.[26] In particular, bandwidth prices are instrumented using marginal costs of electricity generation, a cost shifter. We use panels of electricity generation shares (by coal, gas, oil) at the country-level and quarterly time series data on prices of coal, gas, and oil, as detailed in Section 3.

Table 4 reports the results of the estimation of Equation (19). We show the OLS specification in columns (1) to (3). We find that demand shifters such as country-level broadband subscriptions, GDP, and the number of data centers have a positive effect on bandwidth demand.

The first and second stages of the IV regression are shown in columns (4) and (5) (model IV-1). The first-stage indicates that electricity costs are strong predictors of bandwidth prices.[27] Once the endogeneity of bandwidth prices is accounted for, the price elasticity increases (in absolute value) to -1.36. This is consistent with the expected direction of the bias of the OLS regression. Column (5) also reports the first-stage F-statistic for the weak identification test which indicates that the instrument strongly predicts the endogenous variable.

With respect to the role played by the number of cables, entry of additional cables has a positive and significant effect on demand, holding bandwidth prices fixed. This corresponds to a "market expansion" effect as described in Section 4.1. In Figure 2, we plot the coefficients from column (5) of Table 4. The marginal effect of the number of undersea cables is decreasing, suggesting decreasing marginal returns from diversity.[28] The effect of adding

[25]Controlling for these dynamic trends is necessary in order to accurately estimate the diversification parameters $\{\alpha_n\}$ and avoid over-attribution of demand growth to cable diversity as more cables enter over time.

[26]In our setting, a source of price endogeneity could be the correlation between demand and cost shocks. If marginal costs or the price elasticity depend on the quantity produced, demand shocks may also be transmitted to prices.

[27]We note that bandwidth prices are more correlated with the change in electricity costs. One interpretation is that lagged electricity costs are strongly positively correlated with current bandwidth prices.

[28]We formally test the monotonicity of the effect of the number of cables using likelihood ratio tests for inequality-constrained models (Silvapulle and Sen (2005), Grömping (2010)). We consider two basic test

a second cable (holding prices fixed) on bandwidth demand expands demand as much as a 28.3% decrease in bandwidth prices, adding a third cable is equivalent to a 19.3% decrease in prices, and adding an eighth cable is equivalent to an 7.5% decrease in prices. In Appendix C, we show that the shape and magnitude of these diversity estimates match the predictions of the consumer-level bandwidth provisioning problem, calibrated with data on cable disruptions and repair durations.

In the remainder of the paper, we use estimates based on a specification of $\log(B_{mt})$ as a function of $\log(1 + n_{mt})$ shown in column (6) of Table 4 (model IV-2).[29] The estimated effect of n_{mt} based on this logarithmic specification is plotted as red squares in Figure 2.

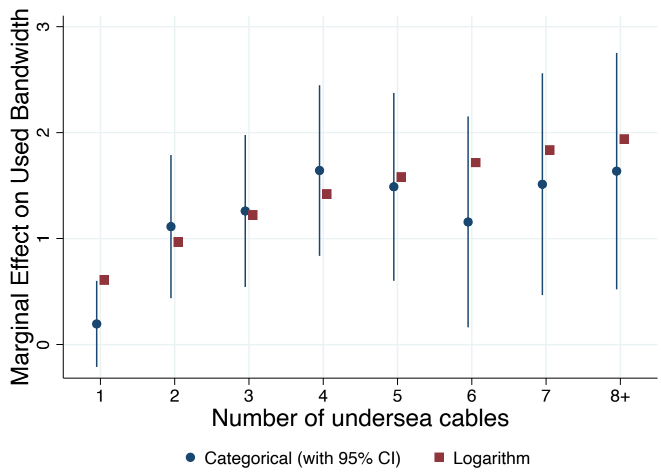

Figure 2: The marginal effect of the number of operating undersea cables on used bandwidth (in log). The specifications with cable count as a categorical variable and with the logarithm of (one plus) the number of cables are shown.

Finally, we conduct robustness checks with respect to effect of the number of cables. As explained in Section 5.1, our baseline specification assumes that firms make entry decisions before the realization of current transient demand shocks: the exact timing of entry, therefore, reflects idiosyncratic shocks (e.g., time-to-build, delays in construction) that are unrelated to unobserved demand shocks. Nonetheless, we address concerns of endogeneity of the number of cables via instrumental variables and a difference-in-difference approach in Appendix E. We find quantitatively similar results with these alternative identification strategies.

problems:
(1) H_0: $\alpha_1 = \alpha_2 = \cdots = \alpha_8$, against H_1: $\alpha_1 \leq \alpha_2 \leq \cdots \leq \alpha_8$ (with at least one strict inequality). The p-value of this test is $6 \cdot 10^{-4}$. Equality of all restrictions is rejected.
(2) H_0: $\alpha_1 \leq \alpha_2 \leq \cdots \leq \alpha_8$, against H_1: at least one inequality violated (e.g., $\alpha_n > \alpha_{n+1}$). The p-value of this test is 0.5484. This null hypothesis is not rejected, i.e., the data do not provide support that these monotonicity restrictions are not all true.

[29]The first-stage in model IV-2 is omitted; estimates are very close to those shown in Column (4).

Table 4: Estimation Results: Used Bandwidth (in log)

	OLS			IV-1		IV-2
	(1)	(2)	(3)	(4) First-stage	(5) Second-stage	(6) Second-stage
Bandwidth Price (10G, log)	-0.926	-0.733	-0.214		-1.360	-1.437
	(0.0876)	(0.0605)	(0.0530)		(0.295)	(0.297)
Number of undersea cables						
One cable	0.388	0.299	0.158	0.0165	0.195	
	(0.669)	(0.802)	(0.167)	(0.0808)	(0.208)	
Two cables	0.631	0.665	1.072	0.0198	1.114	
	(0.666)	(0.816)	(0.338)	(0.110)	(0.345)	
Three cables	1.112	1.072	1.412	-0.135	1.261	
	(0.687)	(0.806)	(0.357)	(0.129)	(0.367)	
Four cables	1.322	1.299	1.835	-0.171	1.643	
	(0.819)	(0.853)	(0.407)	(0.155)	(0.410)	
Five cables	1.413	1.386	1.765	-0.243	1.489	
	(0.731)	(0.842)	(0.434)	(0.192)	(0.452)	
Six cables	1.515	1.416	1.561	-0.363	1.157	
	(0.741)	(0.890)	(0.500)	(0.199)	(0.508)	
Seven cables	2.566	1.507	2.080	-0.507	1.513	
	(0.772)	(0.831)	(0.503)	(0.238)	(0.534)	
Eight or more cables	1.460	1.504	2.305	-0.589	1.637	
	(0.835)	(0.947)	(0.506)	(0.227)	(0.569)	
Demand factors						
Fixed Broadband Subscriptions (log)		0.370	0.208	0.174	0.411	0.420
		(0.116)	(0.111)	(0.0771)	(0.133)	(0.132)
GDP (log)		-0.0646	0.932	0.0221	0.972	1.288
		(0.135)	(0.302)	(0.197)	(0.300)	(0.312)
Aggregate trade flow (log)		-0.00272	-0.0603	0.0261	-0.0317	0.00808
		(0.0854)	(0.0496)	(0.0378)	(0.0770)	(0.0850)
Number of data/cloud centers (log)		0.622	0.165	0.0194	0.193	0.188
		(0.0800)	(0.0800)	(0.0633)	(0.115)	(0.118)
Distance (km, log)		-0.0203				
		(0.134)				
Common official language		0.391				
		(0.214)				
Contiguous		0.694				
		(0.264)				
Electricity price (log)				0.0510		
				(0.0210)		
% change in Electricity price				-0.218		
				(0.0349)		
Number of undersea cables (log)						0.882
						(0.284)
Country Pair FEs	No	No	Yes	Yes	Yes	Yes
Region Pair × Year FEs	No	No	Yes	Yes	Yes	Yes
Weak Identification test					22.31	22.58
Endogeneity test (p-value)					26.8 (0)	28.1 (0)
R^2	0.39	0.71	0.97		0.96	0.95
Adjusted R^2	0.39	0.71	0.97		0.95	0.95
Observations	4908	3863	3849	3863	3863	3863

Note: The unit of observation is a country pair by quarter. Standard errors are clustered at the country pair level. Columns (4) and (5) show the first and second stages of the IV regressions where $\log(P_{mt})$ *is instrumented. Column (6) shows the same second stage as Column (5), but the number of cables is included in logarithm instead of as a categorical variable. Distance corresponds to the bilateral distances between countries, calculated as a weighted arithmetic average of the geodesic distances between the main cities in these countries, where population weights are used. For unilateral variables (GDP, fixed broadband subscriptions, data centers, electricity prices), we restrict the coefficient to be the same for both countries in a pair.*

There are potential limitations to this demand specification. The first set are akin to issues arising in studies of airline markets (with stops and connections). First, bandwidth deployed in a given market (e.g., US–UK) may also include traffic originating from other countries that lack direct cable connections (e.g., traffic between Belgium and the US transiting through the UK), whose characteristics are not controlled for in the regression ($\mathbf{X}_{mt}$). We note, however, that the demand model has a very high fit ($R^2 = 95\%$) indicating that any such unobserved characteristics have limited explanatory power.

Second, we assume that bandwidth deployed in a market (B_{mt}) is carried only by cables *directly* connecting the country pair (n_{mt}), and is not connecting via other countries. Discussions with industry professionals suggest that this indeed tends to be the case because of connection costs and latency considerations. This assumption is reflected in our focus on markets with direct cable connections, i.e., we do not attempt to model the unobserved and complex routing of global data flows. We provide further justifications for these assumptions in Appendix H.

Third, the demand system does not explicitly account for substitution patterns across firms, due to a lack of data on firm-level market shares. This concern is alleviated by the relative homogeneity of bandwidth as a product and the limited degree of supplier differentiation.

Purchased Bandwidth. When computing transfers from consumers to firms ($P_{mt}Q_{mt}$) and firm profits, we convert used bandwidth B_{mt} into purchased bandwidth Q_{mt} using the used-to-purchased bandwidth ratio $\widetilde{f}(n_{mt}, \gamma)$ given by the buyer's first-order condition (Equation (3)). To compute the function $\widetilde{f}(n, \gamma)$, we first estimate the distribution of disruption shocks using data on cable faults and repair durations, then calibrate γ to match the reduced form estimates from Equation (19). Details are provided in Appendix C. In the remainder of the paper, we use a calibrated value of γ equal to 120, but also test the sensitivity of our counterfactual predictions for $\gamma \in [40, 200]$ in Section 6.

Demand State Variable. The estimated demand state variable d_{mt} is defined as the intercept of the (log) demand curve:

$$d_{mt} \equiv \widehat{\alpha}_2 \mathbf{X}_{mt} + \widehat{\eta}_m + \widehat{\eta}_{r(m)t} + \widehat{\epsilon}_{mt} \,. \tag{20}$$

Note that while this equation defines an estimator of d_{mt}, we omit the hat notation to aid legibility in what follows. The residual of the IV regression is included in d_{mt} as it captures omitted time-varying demand shifters.[30] The market and region-year dummies are

[30] Firms make entry decisions after observing d_{mt}, therefore, current period actions a_{it} are correlated with ϵ_{mt}. However, market structure n_{mt} at the beginning of the period is not correlated with ϵ_{mt} as long as demand shocks are not serially correlated, which we verify using the demand residuals. The estimates when instrumenting n_{mt} (Table A3 in Appendix E) also provide support for this assumption.

also included to control for market-specific unobserved heterogeneity and regional trends. Using the demand estimation toward the construction of a demand state variable has the benefit of allowing us to combine several observed and unobserved demand factors affecting bandwidth demand (i.e., the $\mathbf{X}_{mt}$'s and fixed effects) into a single index. This is particularly helpful in anticipation of estimation and computation of equilibria of the dynamic game as it drastically reduces the dimensionality of the state space.

For purpose of illustration, we include plots of used bandwidth B_{mt} in level and logarithm (top and middle panels) and the demand state d_{mt} (bottom panel) for a sample of markets in Figure A8 of Appendix I. Demand for bandwidth increases exponentially over the sample period, with some degree of heterogeneity across markets of different sizes. This exponential growth is the result of both the widespread adoption of the internet on the demand side, with increased usage due to new activities such as social media, streaming, and cloud computing and storage; and on the supply side, advancements in telecommunication technologies which enable faster and more cost-efficient data transmission. The demand state d_{mt} corresponds to the portion of total used bandwidth B_{mt} that is not explained by falling prices (due to cost improvements or competition) or greater diversity but instead is driven by the evolution of exogenous demand factors.

The estimates from Equation (19) imply that the dramatic reduction in prices (the term $\alpha_p \log(P_{mt})$) is the main driver of the observed increase in bandwidth demand. On average, prices decrease by 6.7% annually or by 71% over our timeframe: based on the estimated price elasticity, this accounts for 75% of the increase in bandwidth demand over 2005-2021. In comparison, the increase in the number of cables explains 13% of the growth in bandwidth demand. The remainder is due to exogenous demand factors. We note that falling bandwidth prices likely affect the cost of supplying downstream internet products (e.g., smartphones adoption, cloud computing and storage, or data-intensive applications such as social media), and therefore, may also be reflecting an increase in the quality and variety of downstream internet products.

Transitions of the Demand State. We specify the evolution of the demand state variable d_{mt} as a first-order autoregressive process (AR(1)), where the auto-regressive parameter is homogeneous across markets but the mean varies over markets (i.e., market-specific drift term κ_m), as follows

$$d_{mt} = \kappa_m + \rho_d d_{m,t-1} + \nu_{mt} \,. \tag{21}$$

The parameters of this AR(1) process are estimated by full maximum likelihood.[31] The estimate of the autoregressive coefficient is 0.929 (se = 0.005). Other estimation methods

[31] We have a long panel (16 years of quarterly data) which alleviates concerns related to the incidental parameter problem.

(within group, first-difference, and OLS) yield similar results. We also explored alternative OLS specifications that include a time trend, shown in Appendix Table A9. Column (1) suggests that the process for d_{mt} generally trends upward. Once we include both the time trend and lagged value $d_{m,t-1}$, in column (3), the time trend is no longer statistically significant. This may be due to multicolinearity between the time trend and $d_{m,t-1}$. Additionally, the stationary AR(1) process in Equation (21) can present an upward trend in the short run, because the initial value d_{m0} is below the long-run expectation $\frac{\kappa_m}{1-\rho_d}$.

In the remainder of the estimation, we proceed using the specification in column (2) of Table A9. Appendix H shows the model fit and provides evidence that this specification fits the data well.

In anticipation of the estimation of the dynamic game, we use the method in Tauchen (1986) to discretize this AR(1) process and obtain the transition matrix of the discretized variable on a finite support. The support of the demand state is allowed to vary by market.

5.4 Marginal Costs Estimates

In this section, we present the estimation of cable operators' marginal costs of bandwidth. To estimate these costs, we combine the demand estimates from the previous section and the first-order conditions of the firms' period profit maximization problem. We assume Cournot competition between cable operators that are active in period t. This conduct assumption is motivated by features of the industry and discussion with industry experts: cable operators light capacity in advance and every period (i.e., either activate new fiber pairs or new wavelengths on lit fiber pairs) and are committed to selling their lit capacity. This is in line with strategic behavior under the model of Cournot competition. We test this conduct assumption in Appendix H.[32]

Marginal costs of supplying bandwidth are assumed to be symmetric across suppliers but can vary by market and over time to capture heterogeneity in cost factors across regions and technological advances as discussed below. The symmetry assumption is motivated by our data source which does not report cable-level market shares. We explore the robustness of our results to relaxating the cost symmetry assumption in Appendix F.3.

Recovering Marginal Costs. In market m and period t, firm i's first-order condition with respect to its output q_{imt} is

$$P_{mt}(q_{imt}, q_{-i,mt}) + \frac{\partial P_{mt}}{\partial q_{imt}} q_{imt} = mc_{mt}\,, \tag{22}$$

[32] One implication of this assumption is that a single price P_{mt} prevails for each market-period. We describe our raw price data and how it is aggregated to the market-period level in Section 3.3. While in our raw data some intra-market-period variation in prices exists, it is limited: a regression of quoted prices on market by period fixed effects yields an R^2 value of 90.6%.

where, under a symmetric equilibrium, q_{imt} equals Q_{mt}/n_{mt}.

Equation (22) is used to infer the marginal costs, mc_{mt}, of suppliers active in each market and time period. mc_{mt} corresponds to the marginal cost of supplying one unit of the product (10Gbps of bandwidth over a calendar quarter). Figure 3 shows box-plots of marginal costs estimates for a subsample of region pairs over time.

The downward trend in marginal costs arises from two main factors. First, economies of scale may enable cable operators to expand capacity with decreasing unit costs. Second, technological advances reduce the cost of deploying additional bandwidth. These include fiber-optic transmission upgrades and improvements in software-defined networking that benefit both new entrants and incumbents (see Appendix F.3 (footnote 50) for examples). Significant cost heterogeneity exists across region pairs: the lowest costs are observed in Trans-Atlantic and Intra-Europe markets, while the highest costs are found in Intra-Asia-Oceania and Latin America-North America markets. Within a region pair, costs can vary by an order of magnitude across markets.

We investigate how marginal costs vary with cable characteristics and other cost shifters and the relevance of economies of scale in Appendix G.

Transition of Marginal Cost State. In anticipation of the estimation of the dynamic game, we recover the transition process of marginal costs. We define the cost state variable h_{mt} as the logarithm of the marginal cost, that is,

$$h_{mt} \equiv \log(mc_{mt}) \,, \tag{23}$$

where, as with d_{mt}, we omit the hat notation for legibility. The evolution of the cost state variable h_{mt} is specified as a non-stationary AR(1) process with a time trend, where the auto-regressive parameter is homogeneous across markets but the mean varies over markets (market-specific drift term) as follows,

$$h_{mt} = \tau_m + \rho_h h_{m,t-1} + \delta t + \xi_{mt} \,. \tag{24}$$

The parameters of the AR(1) process are estimated by maximum likelihood. Other estimation approaches (within group, OLS) give similar results. The autoregressive coefficient estimate is 0.887 (se = 0.001) and the time trend δ estimate is -0.024 (se = $6 \cdot 10^{-6}$). As with the demand state variable d_{mt}, we discretize the AR(1) process for h_{mt} to allow for tractable estimation of the dynamic game. Because the process is non-stationary, we obtain the transition matrix for each period. The support of the cost state h_{mt} is allowed to vary by market. Appendix H illustrates model fit and provides evidence that this specification fits the data well.

Figure 3: Marginal Cost Estimates by Region Pair

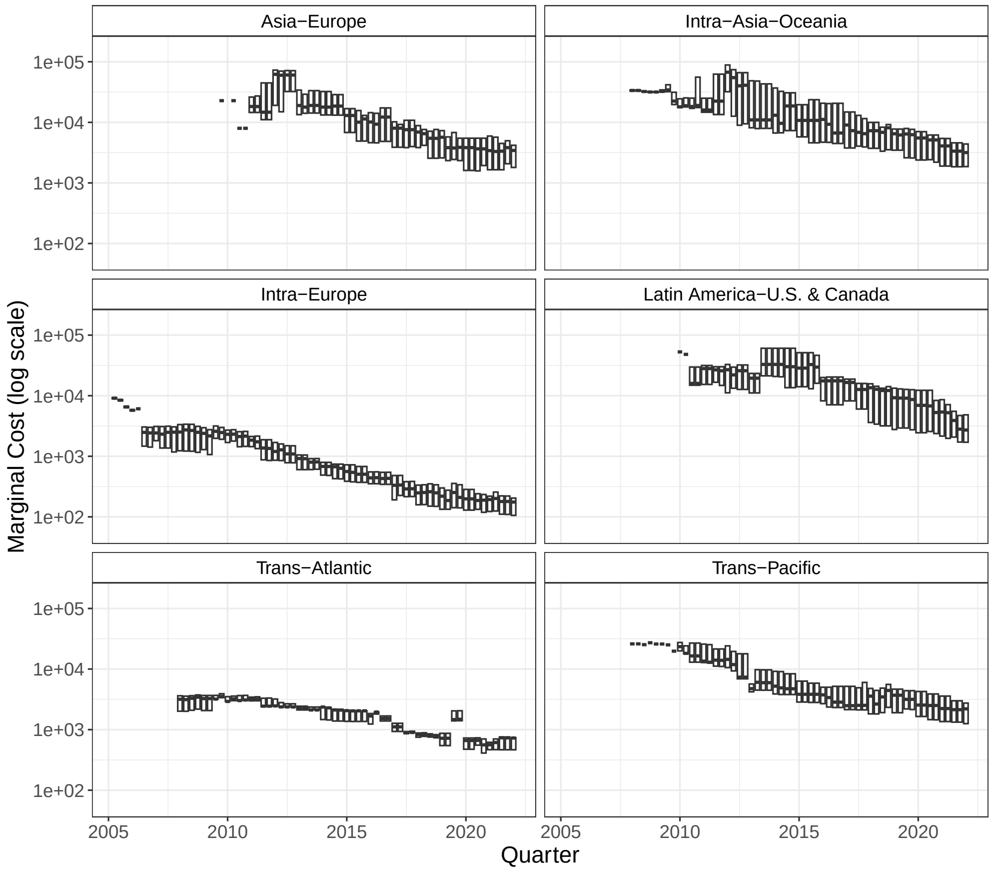

Notes: The figures display box-plots of marginal costs estimates for a subset of region pairs over time. Each white bar shows the first and third quartiles of the marginal costs distribution across markets (country pairs) in a given region pair. The black segment shows the median of this distribution. The unit is US$ per quarter at 10Gbps capacity.

5.5 Dynamic Investment Cost Estimates

In the last step of our estimation approach, we combine the static estimates (from the demand model and firms' marginal costs) to construct period profits and incorporate these profits into firms' dynamic game of entry. The dynamic parameters of interest are the entry costs, fixed costs, and the scale of the private information shocks.

As discussed in the identification of the model (Section 5.1), we cannot separately identify entry costs from fixed costs because we do not observe sufficiently many exit events. Therefore, we start by recovering entry costs directly from data on cable construction costs. Next, we follow the methodology outlined in Section 5.2 to estimate the level of fixed costs and the scale of the logit shocks.

Entry Costs. Our data source provides construction cost information for 47% of active cables. The main cost drivers for undersea cable construction are: (1) undersea components, including fiber, cable, repeaters, and branching units; (2) the "dry plant," encompassing

Submarine Line Termination Equipment, power feeding equipment, monitoring systems, and network protection equipment; and (3) marine operations, such as route surveying, cable loading, laying, and burial. Undersea components are the most distance-sensitive and typically the costliest, as they determine the quantity of cable, fiber, and repeaters required. In our sample, the average construction cost is $244 million, while the median is $66 million. Figure A11 in Appendix I shows cable-level construction costs relative to cable length (in level and log).

We use the construction cost data to predict entry costs for all remaining markets. This is done by regressing construction costs on cable length and region fixed effects, which capture regional differences in topology: e.g., geographical features, such as shallow waters or fault-prone areas, significantly impact fiber installation costs. Entry costs are predicted using the average cable length of all operating cables in a market, as cable length data is available for all active cables. For baseline estimates, we assume entry costs remain constant over time within market. Appendix F.4 relaxes this assumption, modeling time-dependent entry and fixed costs to account for technological advances in cable construction. These advances modestly reduce entry costs over time, and our main conclusions are not sensitive to this more flexible specification.

Fixed Costs and Scale Parameter. Fixed operational costs include administrative costs for office equipment, administrative and network staffing, marketing, legal and regulatory fees, as well as network operation costs for maintenance and upgrade of undersea equipment, landing stations, and network operations facilities (see Appendix F.3 for examples of such system upgrades).

Given the large heterogeneity in cable lengths (e.g., Belgium-United Kingdom and Japan-United States), we allow the fixed costs and the scale of the logit shocks to vary across market. In particular, we parameterize the (quarterly) fixed costs and scale parameter as linear functions of the entry costs: $FC_m = \delta_{FC} EC_m$ and $\theta^{\epsilon}_m = \delta_\epsilon EC_m$.

Although the parameters $(\delta_{FC}, \delta_\epsilon)$ are separately identified because static profits are "observed" when estimating the dynamic game, in practice, we encountered the presence of multiple local maxima of the likelihood function in the ratio $\frac{\delta_{FC}}{\delta_\epsilon}$.[33] To address this problem, we estimate the scale parameter δ_ϵ under different values of δ_{FC} which are informed by discussion with industry professionals and various cable-level cost reports. In particular, we choose upper and lower bounds for δ_{FC} so that our cost estimates are in line with typical industry ratios of operating expenses and capital expenses, as discussed in the model fit Appendix H.

[33] We conjecture that this due to the entry costs and fixed costs term $(-EC_m - \frac{\beta}{1-\beta} FC_m)$ dominating the remaining term (continuation value) that pins down δ_ϵ.

Given a value for δ_{FC}, the scale parameter is estimated precisely and the iterations of our estimation algorithm converge to the same estimate in a short number of iterations. Estimates of the scale parameter are shown in Table 5 for three different estimators. In each bootstrap iteration, the NPL and spectral algorithms converge to the same fixed point of the NPL mapping (regardless of starting values) and, therefore, the corresponding pseudo-likelihood maximizer is virtually identical (up to the fifth digit). Under the assumption that $\delta_{FC} = 0.2\%$, annual fixed costs are in the range of \$1 million (the mean is \$1.36 million and the median is \$1.1 million) in line with estimates reported in TeleGeography (2022). In the counterfactual analysis, we proceed with estimates obtained from the NPL algorithm under the assumption $\delta_{FC} = 0.2\%$.

Table 5: Parameter Estimates from the Dynamic Model

	$\delta_{FC} = 0$	$\delta_{FC} = 0.2\%$	$\delta_{FC} = 0.8\%$
Scale of entry shocks: $\theta^{\epsilon}_m = \delta_\epsilon EC_m$			
Parameter δ_ϵ			
1-PML	0.280 (0.063)	0.284 (0.062)	0.300 (0.059)
NPL Algorithm	0.272 (0.064)	0.277 (0.062)	0.293 (0.059)
Spectral Algorithm	0.272 (0.064)	0.277 (0.062)	0.293 (0.059)

Note: Standard errors, in parenthesis, are obtained by bootstrap (300 replications).

In Appendix H, we evaluate the model fit and provide various test of the model assumptions regarding conduct and the independence of markets.

6 Counterfactual Analysis

6.1 The Role of Supplier Diversification

In the first counterfactual exercise, we assess how quantitatively important supplier diversification is in shaping the evolution of the industry and the total surplus generated.

To implement this counterfactual exercise, we assume buyers cannot diversify their supplier base. That is, in the demand system (Equation (8)), we set n_{mt} to be one regardless of how many cables operate. Entry of new cables intensifies competition (i.e., raises output and reduce prices) but does not benefit buyers through increased supplier diversification. Therefore, the market expansion effect acting through supplier diversification is shut down. In the context of the demand model of Section 4.1, this is equivalent to setting n equal to one in the cost-minization problem (Problem (1)). Given this counterfactual demand model, we solve for the nonstationary MPE of the game using the methodology outlined in Section 5.2.

Table 6 shows the outcomes under the equilibrium played in the data (DGP) and the counterfactual equilibrium (CF) without diversification. Consumer surplus is computed using Equation (5) (scaled by the appropriate market-level heterogeneity d_{mt}, as noted in footnote 16). We also present the expected disruption costs to the buyers (second term in Equation (1)) separately in the table, as these costs are not necessarily a transfer to firms. To compute producer surplus, we convert used bandwidth into purchased bandwidth using the estimated function $\widetilde{f}(n, \gamma)$. Details regarding the estimation of this function are provided in Appendix C.

When buyers cannot diversify, new cable entry does not expand demand as much as in the market outcome. We find that consumer surplus under the counterfactual scenario amounts to 73% of consumer surplus under the DGP: in other words, supplier diversification accounts for 27% of consumer surplus created over the sample period. The reduction in consumer surplus is driven by two changes: first, when buyers cannot diversify, over-provisioning (e.g., the ratio of purchased-to-used bandwidth Q/B in the model of Section 4.1) increases, raising the direct cost of *used* bandwidth; second, the expected disruption costs increase because buyers face higher variance in disruption risk. In particular, the sum of expected disruption costs over our sample period—where the disruption costs are given by the second term in Equation (1)—are 4.7 times greater in the counterfactual scenario than the baseline case. Counterfactual entry probabilities are 12% lower (from a baseline mean entry probability of 4% in the DGP) because entry no longer expands demand via increased diversity.[34] Producer surplus does not vary significantly, however. This is due to the fact that although variable profits are lower under the counterfactual equilibrium, entry costs and private information shocks tend to dominate the continuation value from being an incumbent. Moreover, the reduction in entry rates is associated with increased market power and increased over-provisioning from buyers, which counters the drop in profits due to lower (used bandwidth) demand.

Average total surplus per market over the sample period (in discounted terms) is equal to US$1.11 billion. Preferences for supplier diversification accounts for 11% of total surplus created in the industry over the sample period.

6.2 The Efficiency of Entry

In the second counterfactual exercise, we compare the equilibrium level of entry in the market outcome to the socially optimal level of entry. The objective is to evaluate whether market forces provide entrants with insufficient or excessive entry motives and to compare the resulting level of supplier diversity to that chosen by a social planner.

There are two distortions affecting entry decisions. First, entry may be excessive due to

[34] We focus here on the likelihood of entry rather than the change in the number of cables. The latter may also reflect the fact that we can simulate forward only starting from a date where price and quantity data are available (which in many instances occurs after 2005), see Appendix A.

Table 6: Counterfactual Outcomes Under No Diversification

	DGP		CF		CF/DGP	
	Mean	Median	Mean	Median	Mean	Median
Entry rates						
Number of cables in 2005-Q1 (data)	1.56	1.00				
Number of cables in 2021-Q4 (data)	3.20	3.00				
Expected number of cables in 2021-Q4	3.46	2.83	3.38	2.79	0.98	0.99
Expected probability of entry	0.064	0.028	0.046	0.026	0.88	0.94
Welfare (2005-Q1 to 2021-Q4) in millions of US$						
Consumer Surplus	684.37	123.25	463.52	85.09	0.73	0.71
Expected disruption costs	30.88	5.34	144.55	26.54	4.76	5.69
Producer Surplus	426.63	305.67	428.34	300.41	1.00	1.00
Total Surplus	1111.00	470.50	891.87	446.30	0.89	0.92
Consumer Surplus / Total Surplus	0.36	0.32	0.31	0.25	0.82	0.83

Note: The first two columns show the data generating process (DGP) outcomes under the observed equilibrium. The next two columns (CF) present counterfactual outcomes when diversification is removed. The final two columns show the ratio of counterfactual to DGP outcomes. The "expected probability of entry" is calculated by averaging the expected probability of entry over all periods, based on equilibrium state transitions and choice probabilities. Expected disruption cost corresponds to the second term in Equation (1). Consumer surplus is net of disruption costs.

standard business-stealing motives: i.e., when entry reduces incumbents' output, the private benefit exceeds the social benefit of entry. Second, entry may be insufficient due to a diversity effect: i.e., the marginal entrant contributes to surplus by increasing supplier diversity that they do not (fully) capture as profits because they cannot perfectly price discriminate (Spence (1976), Mankiw and Whinston (1986)). The supplier diversity and business-stealing effects thus push in opposite directions. Whether the business-stealing or diversity effect dominates depends on the shape of the demand curve and the nature of post-entry competition.

We disentangle these two effects in our setting by simulating two counterfactuals. In the planner's benchmark, we solve for the optimal dynamic entry path chosen by a social planner who maximizes total surplus. In the coordination benchmark, we search for the optimal entry path chosen by a planner who maximizes expected producer surplus. In both benchmarks, the planner takes post-entry competition as given, and is thus internalizing these externalities at the margin of entry.[35] In the first scenario, both business-stealing and diversity effects are internalized, whereas in the second scenario, only the business-stealing effect is accounted for by the planner. We note that the solutions to these dynamic optimization problems are greatly simplified by our assumption that only one potential entrant has an opportunity to enter every quarter.[36] This assumption reduces the dimensionality of the state space for the

[35] In the coordination benchmark, the planner does not coordinate output across active firms but only coordinates their entry decisions while taking the oligopolistic post-entry competition as given. The solution to this problem is not necessarily to allow only one firm to enter because producer surplus can be non-monotonic in the number of firms when buyers value diversification.

[36] This constraint is unlikely to bind in practice because there are no observations where two or more cables

social planner to be the same as that of a player in the dynamic game.

To perform this decomposition, we denote by TS_{DGP}, TS_{CB}, and TS_{FB} total surplus under the market outcome, coordination benchmark, and planner benchmark (i.e., first best). The fraction of the welfare gap explained by diversity effects is calculated as $\frac{TS_{FB}-TS_{CB}}{TS_{FB}-TS_{DGP}}$, and the fraction explained by business-stealing is $\frac{TS_{CB}-TS_{DGP}}{TS_{FB}-TS_{DGP}}$.

The results of this counterfactual exercise are shown in Table 7. In the coordination benchmark, the expected number of cables in 2021 is 23% lower compared to the market outcome, producer surplus increases by 31% as business-stealing is internalized. However, consumer surplus decreases by 9%. The expected number of cables at the end of the sample (2021-Q4) in the DGP is distorted downward by diversity effects and upward by business stealing: the magnitude of diversity distortions in terms of number of entrants ranges from 54% to 125% of the business-stealing distortion. For the average market, we find that business-stealing dominates leading to moderately excessive entry rates. Relative to the market outcome, total surplus under the Planner's benchmark (first-best) is on average 10% higher. We find that 53% of this welfare gap is due to diversity effects, whereas 47% is due to business-stealing.

Heterogeneity Analysis. As a final exercise, we examine heterogeneity in the distortions described above at the market level. In particular, we identify the market features and parameters which determine the size of the diversity distortion relative to business-stealing.

First, we regress the change in total surplus (due to each distortion) on various market-level characteristics. We use $(TS_{FB} - TS_{CB})/TS_{DGP}$ and $(TS_{CB} - TS_{DGP})/TS_{DGP}$ as the dependent variables measuring the change in welfare explained by diversity effects and business-stealing, respectively. Figure 4 shows the OLS coefficients from the regression of these dependent variables on entry costs (EC_m), baseline demand level (κ_m) and cost level (τ_m) estimated from the AR(1) processes, mean growth in the exogenous part of demand ($\frac{d_{m,t+1}}{d_{mt}}$), and mean price-cost margin ($p_{mt} - mc_{mt}$). The latter variable is endogenous, but we include it as the correlation can still be informative.

As shown in Figure 4, the size of the welfare gap due to business stealing effects is increasing in entry costs, decreasing in baseline demand, and increasing in the price-cost margin. These findings are consistent with theoretical predictions (Spence (1976), Mankiw and Whinston (1986)) which suggest that the welfare loss is proportional to the level of excess entry costs and profit diversion from incumbents. Business-stealing is more muted if the market size is large, but is larger if incumbents' margins are high (the higher the price-cost margin, the higher is the loss from a reduction in incumbents' quantities).

The welfare gap due to diversity effects is increasing in the price-cost margin; this is con-

enter a market in the same quarter; this assumption is also imposed in estimation.

Table 7: Counterfactual Outcomes Under the Planner and Coordination Benchmarks

	DGP		CF		CF/DGP	
	Mean	Median	Mean	Median	Mean	Median
Social Planner benchmark						
Entry rates						
Number of cables in 2005-Q1 (data)	1.56	1.00				
Number of cables in 2021-Q4 (data)	3.20	3.00				
Expected number of cables in 2021-Q4	3.46	2.83	3.05	2.43	0.86	0.88
Expected probability of entry	0.064	0.028	0.015	0.000	0.273	0.001
Welfare (2005-Q1 to 2021-Q4) in millions of US$						
Consumer Surplus	684.37	123.25	670.75	109.07	0.96	0.98
Expected disruption costs	30.88	5.34	32.98	5.28	1.06	1.00
Producer Surplus	426.63	305.67	511.48	380.81	1.23	1.16
Total Surplus	1111.00	470.50	1182.23	541.70	1.10	1.06
Consumer Surplus / Total Surplus	0.36	0.32	0.33	0.28	0.88	0.92
Coordination benchmark						
Entry rates						
Number of cables in 2005-Q1 (data)	1.56	1.00				
Number of cables in 2021-Q4 (data)	3.20	3.00				
Expected number of cables in 2021-Q4	3.46	2.83	2.78	2.00	0.77	0.80
Expected probability of entry	0.064	0.028	0.005	0.000	0.146	0.000
Welfare (2005-Q1 to 2021-Q4) in millions of US$						
Consumer Surplus	684.37	123.25	602.05	108.83	0.91	0.97
Expected disruption costs	30.88	5.34	47.84	5.24	1.18	1.00
Producer Surplus	426.63	305.67	542.48	402.35	1.31	1.18
Total Surplus	1111.00	470.50	1144.53	526.46	1.08	1.04
Consumer Surplus / Total Surplus	0.36	0.32	0.31	0.26	0.85	0.88

Note: The first two columns correspond to the data generating process (DGP) and show outcomes under the equilibrium played in the data. The following two columns show counterfactual (CF) outcomes under the social planner and coordination benchmarks. The last two columns show the ratio of the counterfactual outcomes to DGP outcomes. The variable "expected probability of entry" is the average over all periods of the expected probability of entry given equilibrium state transitions and choice probabilities. The expected disruption cost is given by the second term in Equation (1). Consumer surplus is net of the disruption costs.

sistent with greater welfare gains from additional entry in markets where firms have market power, which can be correlated with a limited number of suppliers. The size of the welfare gap is also increasing in baseline demand, since infra-marginal buyers also benefit from entry via greater diversity. Taken together, these findings suggest that policy interventions to *encourage* entry would best target markets with low entry costs, low baseline costs (τ_m), high price-cost margins, and high market size (κ_m). In such markets, distortions due to diversity effects will tend to dominate, leading to socially insufficient entry.

Second, we conduct a comparative statics exercise evaluating the sensitivity of the model's counterfactual predictions to buyers' preferences for diversity. To do this, we vary the disruption cost parameter γ (set to 120 in the baseline estimation) from a value of 40 to 200 (see Appendix C for a derivation of these bounds). For each value of the disruption cost parameter, which affects the degree of diversification needs, we solve for the market out-

Figure 4: Correlation Between the Welfare Gaps and Market Characteristics

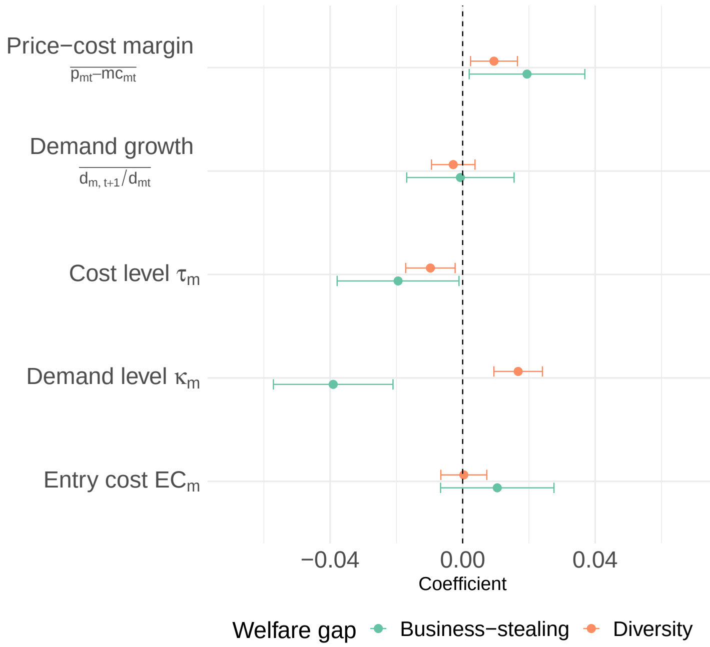

Notes: The figure displays the coefficients from a regression of $(TS_{CB} - TS_{DGP})/TS_{DGP}$ (business-stealing) and $(TS_{FB} - TS_{CB})/TS_{DGP}$ (diversity) on market characteristics shown on the y-axis. 95% confidence intervals are represented as horizontal lines. All variables used as regressors are standardized. For the price-cost margin and demand growth, we take the average of their values in the DGP over all periods.

come, coordination and planner benchmarks, then compute welfare. The relative sizes of the diversity and business-stealing distortions depend crucially on the nature of demand for diversity: we find that under the lower value of disruption costs, the majority of the welfare gap (92%) is due to business-stealing, whereas with the higher disruption costs of 200, 62% of the welfare gap is due to diversity effects.

The takeaways from these two exercises are that (1) suppliers' inability to internalize positive externalities from risk diversification can significantly distort entry decisions, and (2) substantial heterogeneity exists, even within a single industry, in the magnitude of entry distortion caused by supplier diversification. The first exercise shows that factors such as market size, price-cost margin, and baseline cost levels are statistically significant determinants of the welfare gap caused by diversity effects. The second exercise highlights that disruption costs have a first-order influence on the size of the distortion. Whether a market achieves sufficient supplier diversity—and might justify entry subsidies from a social cost-benefit perspective—depends on market- and industry-specific attributes. Policymakers addressing these distortions must consider market-specific characteristics and the magnitude of disruption costs to buyers.

We also note that, while policy interventions would target the entry margin in this unregulated industry, these interventions might take different forms in other contexts. For instance, in natural monopoly industries such as electricity transmission and distribution, disruptions can be mitigated by capital investments, such as the maintenance of aging infrastructure, which are incentivized by regulated rates and auditing (Lim and Yurukoglu (2018)). Ensuring efficient levels of reliability is more challenging in that context due to asymmetric information between the regulator and the firm, compared to our setting where reliability is supplied via the market mechanism.

7 Conclusion

Supply disruptions and the ensuing diversification sought by buyers are critical features in many industries. This paper quantitatively assesses their impact on an industry's dynamic evolution, focusing on the global internet backbone. To achieve this, we build a demand model of bandwidth provisioning and diversification under disruptions, coupled with a dynamic oligopoly game of entry by cable operators. The model is estimated using novel data on used bandwidth, bandwidth prices, cable characteristics, and cable disruptions.

We find that supplier diversification accounts for a significant portion of entry rates and surplus created during the sample period, 2005-2021. Furthermore, we quantify entry bias arising from entrants' inability to capture the benefits of diversity. Specifically, when buyers diversify their supplier base, a marginal entrant increases surplus by enhancing diversity but does not fully capture this gain in profits. In this industry, distortions due to diversity effects are significant and comparable in magnitude to distortions from business-stealing.

Our framework is applied to the global internet backbone, an industry marked by frequent and costly supply disruptions and a strong buyer preference for diverse physical paths. Studying this industry is valuable in its own right, as it provides insights into the infrastructure and physical connections underpinning the internet, a cornerstone of modern communication and commerce.

Our model highlights that in industries with supplier disruptions, supplier diversity has a public good-like quality and may be under-provided by market forces—an under-studied distortion warranting attention, particularly where disruptions impose significant economic costs. The proposed empirical framework could also be applied to other industries facing supply bottlenecks (e.g., energy transportation, essential infrastructure, semiconductor manufacturing) or to evaluate mergers in such sectors. Understanding the role of these incentives in merger analysis and exploring the provision of diversification under monopoly or multi-cable ownership are promising avenues for future research.

References

Ackerberg, Daniel A., and Marc Rysman. 2005. "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects." *The RAND Journal of Economics* 36 (4): 771–788, http://www.jstor.org/stable/4135256.

Aguirregabiria, Victor, and Mathieu Marcoux. 2021. "Imposing equilibrium restrictions in the estimation of dynamic discrete games." *Quantitative Economics*.

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential estimation of dynamic discrete games." *Econometrica* 75 (1): 1–53.

Aguirregabiria, Victor, and Junichi Suzuki. 2014. "Identification and counterfactuals in dynamic models of market entry and exit." *Quantitative Marketing and Economics* 12 (3): 267–304.

APEC. 2012. "Economic Impact of Submarine Cable Disruptions." *Asia-Pacific Economic Cooperation Policy Support Unit*, https://www.apec.org/publications/2013/02/economic-impact-of-submarine-cable-disruptions.

Arcidiacono, Peter, and Robert A Miller. 2011. "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity." *Econometrica* 79 (6): 1823–1867.

Barrot, Jean Noël, and Julien Sauvagnat. 2016. "Input specificity and the propagation of idiosyncratic shocks in production networks." *Quarterly Journal of Economics* 131 1543–1592.

Berry, Steven, Alon Eizenberg, and Joel Waldfogel. 2016. "Optimal product variety in radio markets." *The RAND Journal of Economics* 47 (3): 463–497.

Besen, Stanley, Paul Milgrom, Bridger Mitchell, and Padmanabhan Srinagesh. 2001. "Advances in Routing Technologies and Internet Peering Agreements." *American Economic Review* 91 (2): 292–296.

Blum, Bernardo S., and Avi Goldfarb. 2006. "Does the internet defy the law of gravity?" *Journal of International Economics* 70 384–405.

Caillaud, Bernard, and Bruno Jullien. 2003. "Chicken Egg: Competition among Intermediation Service Providers." *The RAND Journal of Economics* 34 309–328.

Caoui, El Hadi, and Andrew Steck. 2024. "Content Providers and the Deployment of Internet Infrastructure."

Carvalho, Vasco M., Makoto Nirei, Yukiko U. Saito, and Alireza Tahbaz-Salehi. 2021. "Supply chain disruptions: Evidence from the great east japan earthquake." *Quarterly Journal of Economics* 136 1255–1321.

Castro-Vincenzi, Juanma. 2022. "Climate Hazards and Resilience in the Global Car Industry."

Collard-Wexler, Allan. 2013. "Demand Fluctuations in the Ready-Mix Concrete Industry." *Econometrica* 81 1003–1037.

Conte, Maddalena, Pierre Cotterlaz, and Thierry Mayer. 2021. "The CEPII gravity database." *CEPII: Paris, France*.

Crew, M. A., and P. R. Kleindorfer. 1978. "Reliability and public utility pricing." *American Economic Review* 68 (1): 31–40.

Crémer, Jacques, Patrick Rey, and Jean Tirole. 2000. "Connectivity in the Commercial Internet." *The Journal of Industrial Economics* 48 433–472.

Draganska, Michaela, Michael Mazzeo, and Katja Seim. 2009. "Beyond plain vanilla: Modeling joint product assortment and pricing decisions." *Quantitative Marketing and Economics* 7 105–146.

Dzieza, Josh. 2024. "The Cloud Under the Sea." April, theverge.com [Online; posted 16-April-2024].

Eizenberg, Alon. 2014. "Upstream innovation and product variety in the U.S. home PC market." *Review of Economic Studies* 81 1003–1045.

Elliott, Matthew, Benjamin Golub, and Matthew V. Leduc. 2022. "Supply Network Formation and Fragility." *American Economic Review* 112 2701–2747.

Fan, Ying, and Chenyu Yang. 2020. "Competition, product proliferation, and welfare: A study of the US smartphone market." *American Economic Journal: Microeconomics* 12 99–134.

Farrell, Joseph, and Nancy T. Gallini. 1988. "Second-sourcing as a commitment: Monopoly incentives to attract competition." *Quarterly Journal of Economics* 103 (4): 673–694.

Federal Communications Commission. 2015. "FCC Record 15-119: A Comprehensive Compilation of Decisions, Reports, Public Notices, and Other Documents."

Garbade, Kenneth D., and William L. Silber. 1978. "Technology, Communication, and the Performance of Financial Markets: 1840-1975." *The Journal of Finance* 33 819–832.

Greenstein, Shane. 2015. *How the internet became commercial: Innovation, privatization, and the birth of a new network.* Princeton University Press.

Grossman, Gene M., Elhanan Helpman, and Hugo Lhuillier. 2023. "Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring?" *Journal of Political Economy.*

Grömping, Ulrike. 2010. "Inference with Linear Equality and Inequality Constraints Using R: The Package ic.infer." *Journal of Statistical Software* 33 (10): 1–31.

Hjort, Jonas, and Jonas Poulsen. 2019. "The arrival of fast internet and employment in Africa †." *American Economic Review* 109 1032–1079.

Hoag, Christopher. 2006. "The Atlantic Telegraph Cable and Capital Market Information Flows." *The Journal of Economic History* 66 342–353.

Igami, Mitsuru. 2017. "Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998." *Journal of Political Economy* 125 (3): 798–847.

Jeon, Jihye. 2022. "Learning and investment under demand uncertainty in container shipping." *The RAND Journal of Economics* 53 (1): 226–259.

Jeon, Jihye, and Marc Rysman. 2024. "International Trade in Data: Investment and Usage of the Subsea Internet Cable Network."

Joskow, Paul, and Jean Tirole. 2007. "Reliability and competitive electricity markets." *RAND Journal of Economics* 38 (1): 60–84.

Judd, Kenneth L. 1998. *Numerical Methods in Economics.* Volume 1. of MIT Press Books, The MIT Press.

Kalouptsidi, Myrto. 2014. "Time to Build and Fluctuations in Bulk Shipping." *American Economic Review* 104 564–608.

Laffont, Jean-Jacques, Scott Marcus, Patrick Rey, and Jean Tirole. 2001. "Internet Peering." *The American Economic Review* 91 287–291.

Lim, Claire S.H., and Ali Yurukoglu. 2018. "Dynamic natural monopoly regulation: Time inconsistency, Moral hazard, and political environments." *Journal of Political Economy* 126 (1): 263–312.

Malphrus, Steve. 2009. "Keynote Address, Reliability of Global Undersea Communication Cables Infrastructure (ROGUCCI), Summit Dubai, UAE."

Mankiw, N Gregory, and Michael D Whinston. 1986. "Free Entry and Social Inefficiency." *The RAND Journal of Economics* 17 48–58.

Pakes, Ariel. 2017. "Empirical Tools and Competition Analysis: Past Progress and Current Problems." *International Journal of Industrial Organization* 241–266.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2008. "Asymptotic least squares estimators for dynamic games." *Review of Economic Studies* 75 (3): 901–928.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2010. "Sequential Estimation of Dynamic Discrete Games: A Comment." *Econometrica* 78 (2): 833–842.

Reverdy, Dominique, and Ivan Skenderoski. 2015. "Submarine Cables: Structuring and Financing Options." *Salience Whitepaper*.

Rust, John. 2000. "Parametric Policy Iteration : An Efficient Algorithm for Solving Multidimensional DP Problems?."Technical report.

Seixas, Pedro. 2015. "Broadband costing and pricing - Approaches and best practices." *International Telecommunication Union*.

Shepard, Andrea. 1987. "Licensing to Enhance Demand for New Technologies." *The RAND Journal of Economics* 18 (3): 360.

Silvapulle, Mervyn J., and Pranab Kumar Sen. 2005. *Constrained statistical inference: Inequality, order and shape restrictions*. United States of America: John Wiley Sons, , 1st edition.

Spence, Michael. 1976. "Product Selection, Fixed Costs, and Monopolistic Competition." *The Review of Economic Studies* 43 (2): 217–235.

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

Sweeting, Andrew. 2013. "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry." *Econometrica* 81 (5): 1763–1803.

Swinhoe, Dan. 2022a. "The cable ship capacity crunch." December, datacenterdynamics.com [Online; posted 6-December-2022].

Swinhoe, Dan. 2022b. "Cable ship shortage sees Maple Leaf Fibre drop plans for fiber cable under Lake Ontario." May, datacenterdynamics.com [Online; posted 17-May-2022].

Tauchen, George. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economics Letters* 20 177–181.

TeleGeography. 2022. "Product Documentation." teleogeography.com .

The Guardian. 2016. "Buried at sea: the companies cashing in on abandoned cables." December, theguardian.com [Online; posted 14-December-2016].

Wollmann, Thomas G. 2018. "Trucks without bailouts: Equilibrium product characteristics for commercial vehicles." *American Economic Review* 108 1364–1406.

Online Appendix

A Details of the Market Definition

This section provides details on the markets contained in the TeleGeography datasets that are used in estimation. As discussed in Section 3.3, we define a market as a country-pair. Our estimation panel is not balanced; new cables are continually built, connecting new country-pairs resulting in new markets appearing in the data. This is illustrated in the left panel of Figure A1.

Figure A1: Cumulative Number of Markets (left) and Cumulative Bandwidth Coverage (right) in the Estimation Panel

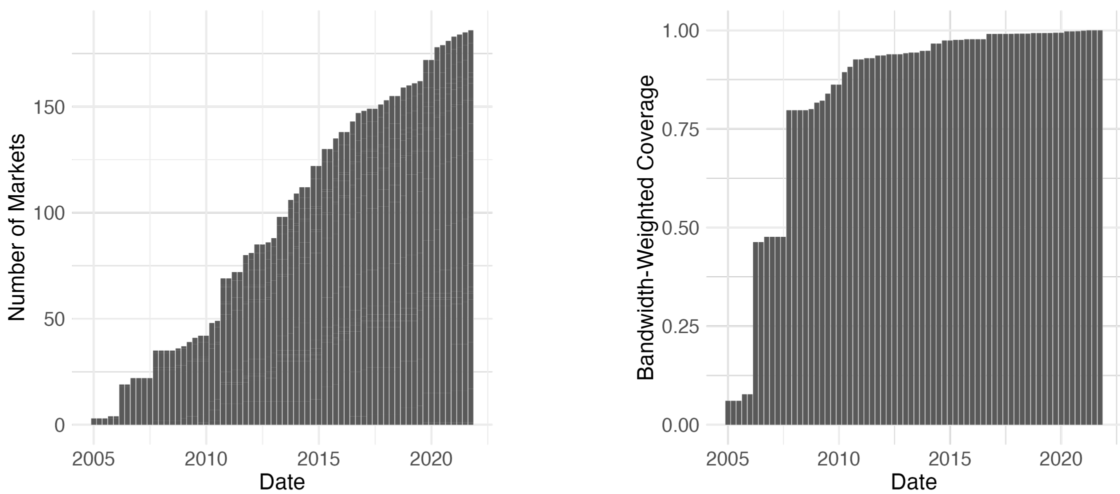

Notes: The left figure shows the cumulative number of distinct country-to-country markets for which both price and used bandwidth data are available. The right figure shows final-period bandwidth-weighted cumulative coverage in the estimation panel.

However, on a bandwidth-weighted basis, the estimation panel is well balanced, as is shown in Figure A1 (right panel). The right panel of Figure A1 plots the final-period bandwidth-weighted cumulative coverage of the estimation panel (final-period bandwidth is used to account for overall bandwidth growth over the data). It can be seen that beginning in the fourth quarter of 2007 the panel's coverage reached 80% of aggregate final-period bandwidth and thereafter remained substantially complete.

We provide further details on the markets included in our sample in Tables A1 and A2. Table A1 contains information on the largest 15 markets (by bandwidth) in the estimation panel, showing the first and last period for which information (on prices and bandwidth used) is available in the data.

Table A2 shows similar data broken down by region pair. For the sake of brevity, only the largest three markets in each region pair (by bandwidth) are shown.

Table A1: Top 15 Markets

Market	First Period	Final Period	Observations	Bandwidth
Germany-United Kingdom	2006-Q2	2021-Q4	63	133,950.7
France-Germany	2006-Q2	2021-Q4	63	131,869.1
Germany-Netherlands	2006-Q2	2021-Q4	63	114,501.2
United Kingdom-United States	2007-Q4	2021-Q4	57	110,773.3
France-United Kingdom	2006-Q2	2021-Q4	63	103,257.7
Netherlands-United Kingdom	2006-Q2	2021-Q4	63	98,091.4
France-Spain	2006-Q2	2021-Q4	63	96,661.1
Japan-United States	2007-Q4	2021-Q4	57	93,189.8
Denmark-United States	2007-Q4	2021-Q4	12	91,767.4
Ireland-United Kingdom	2007-Q4	2021-Q4	57	85,333.5
Spain-United States	2007-Q4	2021-Q4	20	71,005.0
Denmark-Sweden	2005-Q1	2021-Q4	68	61,609.9
Denmark-Germany	2005-Q1	2021-Q4	37	59,898.3
France-United States	2007-Q4	2021-Q4	57	58,693.1
Ireland-United States	2007-Q4	2021-Q4	55	53,448.6

Note: This table presents details on the top 15 markets in the data, in terms of used bandwidth in the final period. Columns 2 and 3 contain the first and last period each market appears in our estimation panel. Column 4 contains the number of periods for each market in the final estimation panel. Column 5 contains the used bandwidth in the final period in Gbps.

Table A2: Top 3 Markets by Region Pair

Market	First Period	Final Period	Bandwidth	Market	First Period	Final Period	Bandwidth
Africa-Africa				**Europe-Europe**			
Kenya-South Africa	2018-Q4	2021-Q4	827.4	Germany-U.K.	2006-Q2	2021-Q4	133,950.7
Ghana-Nigeria	2017-Q2	2021-Q4	223.6	France-Germany	2006-Q2	2021-Q4	131,869.1
Angola-South Africa	2019-Q4	2021-Q3	204.4	Germany-Netherlands	2006-Q2	2021-Q4	114,501.2
Africa-Asia				**Europe-Middle East**			
Egypt-India	2013-Q4	2015-Q1	23.1	France-U.A.E.	2011-Q4	2021-Q4	3,516.6
India-Kenya	2013-Q4	2019-Q4	11.0	France-Saudi Arabia	2011-Q4	2021-Q4	3,366.5
India-Tanzania	2013-Q4	2015-Q4	0.7	France-Oman	2014-Q4	2021-Q4	2,673.1
Africa-Europe				**Europe-U.S. & Canada**			
Egypt-France	2014-Q4	2021-Q4	5,885.6	U.K.-U.S.	2007-Q4	2021-Q4	110,773.3
Nigeria-U.K.	2015-Q2	2021-Q4	2,596.2	Denmark-U.S.	2007-Q4	2021-Q4	91,767.4
France-South Africa	2018-Q4	2021-Q1	1,911.4	Spain-U.S.	2007-Q4	2021-Q4	71,005.0
Africa-Latin America				**Latin America-Latin America**			
Angola-Brazil	2019-Q4	2021-Q4	164.0	Argentina-Chile	2010-Q2	2021-Q4	6,546.6
				Argentina-Brazil	2010-Q2	2021-Q4	5,957.5
				Chile-Peru	2010-Q4	2021-Q4	2,211.2
Africa-Middle East				**Latin America-U.S. & Canada**			
Kenya-U.A.E.	2018-Q1	2020-Q3	282.5	Brazil-U.S.	2009-Q4	2021-Q4	45,624.8
Egypt-Oman	2013-Q4	2015-Q1	0.2	Mexico-U.S.	2010-Q2	2021-Q4	29,696.1
				Chile-U.S.	2010-Q2	2021-Q4	11,423.4
Asia-Asia				**Middle East-Middle East**			
Japan-Singapore	2009-Q1	2021-Q4	32,841.5	Bahrain-U.A.E.	2014-Q1	2021-Q2	361.5
India-Singapore	2010-Q3	2021-Q4	27,004.4	Kuwait-U.A.E.	2014-Q4	2021-Q2	178.3
Indonesia-Singapore	2014-Q2	2021-Q4	24,756.7	Qatar-U.A.E.	2014-Q1	2021-Q2	154.5
Asia-Europe				**Oceania-Oceania**			
France-India	2016-Q4	2021-Q4	22,429.3	Australia-New Zealand	2013-Q1	2021-Q4	3,827.7
France-Singapore	2011-Q2	2021-Q4	4,011.1				
India-U.K.	2010-Q4	2021-Q4	3,674.8				
Asia-Middle East				**Oceania-U.S. & Canada**			
Saudi Arabia-Singapore	2012-Q2	2021-Q4	1,074.8	Australia-U.S.	2010-Q4	2021-Q4	12,133.2
Singapore-U.A.E.	2012-Q2	2021-Q4	999.8	New Zealand-U.S.	2013-Q1	2021-Q4	1,313.5
India-U.A.E.	2010-Q4	2021-Q4	567.8				
Asia-Oceania				**U.S. & Canada-U.S. & Canada**			
Australia-Singapore	2011-Q4	2021-Q4	3590.6	Canada-U.S.	2005-Q4	2021-Q4	32,980.7
Australia-Japan	2011-Q2	2021-Q4	1933.4				
Australia-Indonesia	2016-Q1	2018-Q2	5.6				
Asia-U.S. & Canada							
Japan-U.S.	2007-Q4	2021-Q4	93,189.8				
Singapore-U.S.	2009-Q3	2021-Q4	34,663.6				
South Korea-U.S.	2012-Q2	2021-Q4	3,774.9				

Note: This table presents details on the top 3 markets by region pair, in terms of used bandwidth in the final period. Columns 2 and 3 contain the first and last period each market appears in our estimation panel. Column 4 contains the used bandwidth in the final period in Gbps.

B Proofs

Lemma 1. *Proof* The integral in Equation (3) is continuous and strictly increasing in $\frac{B}{Q}$, converges to 0 for $\frac{B}{Q}$ close to zero, and equals μ when $\frac{B}{Q}$ is one. A unique interior solution exists by the Intermediate Value theorem. The Implicit Function theorem implies that:

$$\frac{\partial \widetilde{f}(n,\gamma)}{\partial \gamma} = \frac{-1}{\gamma^2 \widetilde{f} g_n(\widetilde{f})} < 0\,,$$

proving that $\widetilde{f}(n,\gamma)$ is decreasing in γ.□

Proposition 1. *Proof* The function $\widetilde{h}(n,\gamma)$ can be expressed as:

$$\begin{aligned}\widetilde{h}(n,\gamma) &= \frac{1}{\widetilde{f}(n,\gamma)} + \gamma\,\mathbb{E}\left[\max\left\{1 - \frac{1}{\widetilde{f}(n,\gamma)}\frac{1}{n}\sum_{i=1}^{n}\delta_i, 0\right\}\right] \\ &= \frac{1}{\widetilde{f}(n,\gamma)} + \gamma\int_0^{\widetilde{f}(n,\gamma)}\left(1 - \frac{u}{\widetilde{f}(n,\gamma)}\right) g_n(u)du \\ &= \frac{1}{\widetilde{f}(n,\gamma)} + \gamma\left(G_n(\widetilde{f}(n,\gamma)) - \frac{1}{\gamma\widetilde{f}(n,\gamma)}\right) \\ &= \gamma G_n(\widetilde{f}(n,\gamma))\,,\end{aligned} \tag{25}$$

where the third equality is obtained by using the first-order condition (3) at $\widetilde{f}(n,\gamma)$. We show that the sequence $G_n(\widetilde{f}(n,\gamma))$ is decreasing. The first-order condition (3) can be rewritten, with the change of variables $p = G_n(u)$, as:

$$\frac{1}{\gamma} = \int_0^{G_n(\widetilde{f}(n,\gamma))} G_n^{-1}(p)dp\,. \tag{26}$$

We distinguish between two possible cases. If $\widetilde{f}(n,\gamma) \leq x_n$, then $G_{n+1}(x) < G_n(x)$ for all $x < \widetilde{f}(n,\gamma)$. Equivalently, $G_n^{-1}(p) < G_{n+1}^{-1}(p)$ for all $p < G_n(\widetilde{f}(n,\gamma))$. Because Equation (26) holds for all n, we have:

$$\frac{1}{\gamma} = \int_0^{G_{n+1}(\widetilde{f}(n+1,\gamma))} G_{n+1}^{-1}(p)dp = \int_0^{G_n(\widetilde{f}(n,\gamma))} G_n^{-1}(p)dp < \int_0^{G_n(\widetilde{f}(n,\gamma))} G_{n+1}^{-1}(p)dp.$$

The inequality implies that $G_{n+1}(\widetilde{f}(n+1,\gamma)) < G_n(\widetilde{f}(n,\gamma))$.

If $\widetilde{f}(n,\gamma) > x_n$, then $G_n^{-1}(p) > G_{n+1}^{-1}(p)$ for all $p \geq G_n(\widetilde{f}(n,\gamma))$. Note that Equation (26) can be rewritten as

$$\frac{1}{\gamma} = \mu - \int_{\widetilde{f}(n,\gamma)}^1 u g_n(u)du = \mu - \int_{G_n(\widetilde{f}(n,\gamma))}^1 G_n^{-1}(p)dp\,.$$

Using a similar argument as in the first case, we can show that:

$$\mu - \frac{1}{\gamma} = \int_{G_{n+1}(\widetilde{f}(n+1,\gamma))}^{1} G_{n+1}^{-1}(p)dp = \int_{G_n(\widetilde{f}(n,\gamma))}^{1} G_n^{-1}(p)dp > \int_{G_n(\widetilde{f}(n,\gamma))}^{1} G_{n+1}^{-1}(p)dp\,.$$

The inequality implies that $G_{n+1}(\widetilde{f}(n+1,\gamma)) < G_n(\widetilde{f}(n,\gamma))$. This proves that $\widetilde{h}(n,\gamma)$ is strictly decreasing in n.

To prove convergence, define the sequence of functions $m_n(x) = \int_0^x G_n^{-1}(p)dp$. This sequence converges point-wise to $m(x) = \mu x$. Because $m_n(x)$ and $m_n'(x) = G_n^{-1}(x)$ are both uniformly bounded by one, m_n converges to m uniformly over $[0,1]$ (Arzelà–Ascoli theorem). In particular,

$$\forall \epsilon > 0, \exists N, \text{s.t.} \forall n \geq N, \left| m_n\left(G_n(\widetilde{f}(n,\gamma))\right) - m\left(G_n(\widetilde{f}(n,\gamma))\right)\right| < \epsilon\,.$$

Since $m_n\left(G_n(\widetilde{f}(n,\gamma))\right)$ equals $\frac{1}{\gamma}$ for all n, the previous statement implies that $m\left(G_n(\widetilde{f}(n,\gamma))\right)$ (i.e., $\mu G_n(\widetilde{f}(n,\gamma)))$ converges to $\frac{1}{\gamma}$, and hence, $\widetilde{h}(n,\gamma) = \gamma G_n(\widetilde{f}(n,\gamma))$ converges to $\frac{1}{\mu}$.□

C Parameterization of the Consumer-Based Bandwidth Demand Model

We calibrate the distribution of disruption shocks δ_i using data on cable faults and the duration of repairs, defined as the duration (in days) between the date the disruption starts and the date at which repairs are completed. We then compare the predictions of the bandwidth provisioning problem to the data.

The disruption shock takes the form $\delta_i = (1 - X_i) + X_i Y_i$, where X_i is a Bernoulli random variable with parameter p_d and $Y_i \in [0,1]$ follows a Beta distribution $Beta(a_1, a_2)$. We assume that X_i and Y_i are independent. That is, a disruption occurs with probability p_d; conditional on a disruption occurring ($X_i = 1$), the duration of repairs is stochastic: the random variable $Y_i \in [0,1]$ corresponds to uptime (per quarter). For example, conditional on a disruption occurring, if Y_i equals 0.5, the cable is operational for half of a quarter (or 1.5 months).[37]

We set the per-quarter probability of disruption p_d between 5% and 10% to match aggregate figures of 100 to 200 cable disruptions per year, out of a total of 434 cables at the end of our sample. Conditional on a disruption occurring, we use the number of repair days available for each cable disruption in our cable fault data to estimate the probability distribution of uptime Y_i. The mean repair duration is 25 days, and the median is 14 days. Panel (a) of Figure A2 shows a histogram of Y_i, constructed using the repair days data.

[37] The variable Y_i equals $\frac{91.25-\text{repair days}}{91.25}$. We cap the repair days at 91.25 days (a quarter) for simplicity. Five cable disruption events (out of 168 in our data) take longer than 91 days to repair.

We estimate the parameters of the Beta distribution by Maximum Likelihood, and plot the corresponding density of $Y_i \sim Beta(\widehat{a}_1, \widehat{a}_2)$ along with the histogram. Under this parameterization, the expectation of δ_i, denoted μ, equals 0.96 if $p_d = 10\%$ or 0.98 if $p_d = 5\%$. The latter corresponds to a 2% expected downtime.

Recall that the logarithm of used bandwidth can be expressed as a function of the bandwidth price and the number of cables as follows (Equation (7))

$$\log(B) = \frac{1}{\alpha}\log(P) + \frac{1}{\alpha}\log\left(\widetilde{h}(n,\gamma)\right). \tag{27}$$

Panel (b) of Figure A2 plots $\frac{1}{\alpha}\log(\widetilde{h}(n,\gamma))$, corresponding to the effect of n on $\log(B)$ predicted by the consumer-level demand model for different values of γ, alongside the reduced-form estimates $\{\widehat{\alpha}_n\}_{n=1}^{8}$ (and confidence bands) from Equations (8).[38] The functions $\widetilde{f}(n,\gamma)$ and $\widetilde{h}(n,\gamma)$ do not have closed-form expressions, but can be easily simulated given the distribution of δ_i. Informed by the mapping between reduced form and structural parameters, we set $\frac{1}{\alpha}$ to the estimated price elasticity ($\widehat{\alpha}_p = -1.36$).

The consumer-level demand model predicts an effect of the number of suppliers n on (log) used bandwidth that is broadly consistent with the reduced form coefficients. In terms of magnitudes, a value of γ in the range $[40, 200]$ best matches our reduced form estimates. We verify that (log) *purchased* bandwidth Q, which is obtained as $\log(Q) = \log(B) - \log(\widetilde{f}(n,\gamma))$, is also increasing and concave in n.

While it is difficult to obtain precise estimates of the economic costs of cable disruptions γ, industry reports estimate damages in the same order of magnitude: for example, APEC (2012) provides simulations of the economic costs of cable disruptions based on estimates of the contribution of the internet economy to GDP and trade, and reliance on international bandwidth.[39] For the year 2012, they calculate economic costs to Australia on the order of US$ 4.6 million per Gbps per month disrupted; on the other hand, the average bandwidth price connecting to Australia in 2012 is US$ 43 thousand per Gbps per month, which results in estimates of γ on the order of 106.[40] We stress that the cost of disruption in our model may capture other less tangible costs (e.g., reputation, backlog) from the perspective of the buyer, which are not accounted for by studies using aggregate statistics such as APEC (2012).

[38] To make the estimates comparable, when plotting $\frac{1}{\alpha}\log\left(\widetilde{h}(n,\gamma)\right)$ as a function of n for different values of γ, we normalize the value at $n = 1$ to match our reduced form estimate ($\widehat{\alpha}_1 = 0.195$). The reduced-form coefficients are measured relative to the omitted category (no cables).

[39] The report and model simulations are available at: https://www.apec.org/publications/2013/02/economic-impact-of-submarine-cable-disruptions [Accessed: September 23rd, 2024].

[40] Similar magnitudes of economic damages are computed for other countries: e.g., US$ 601 thousand for Singapore, US$ 3.8 million for South Korea per Gbps per month disrupted.

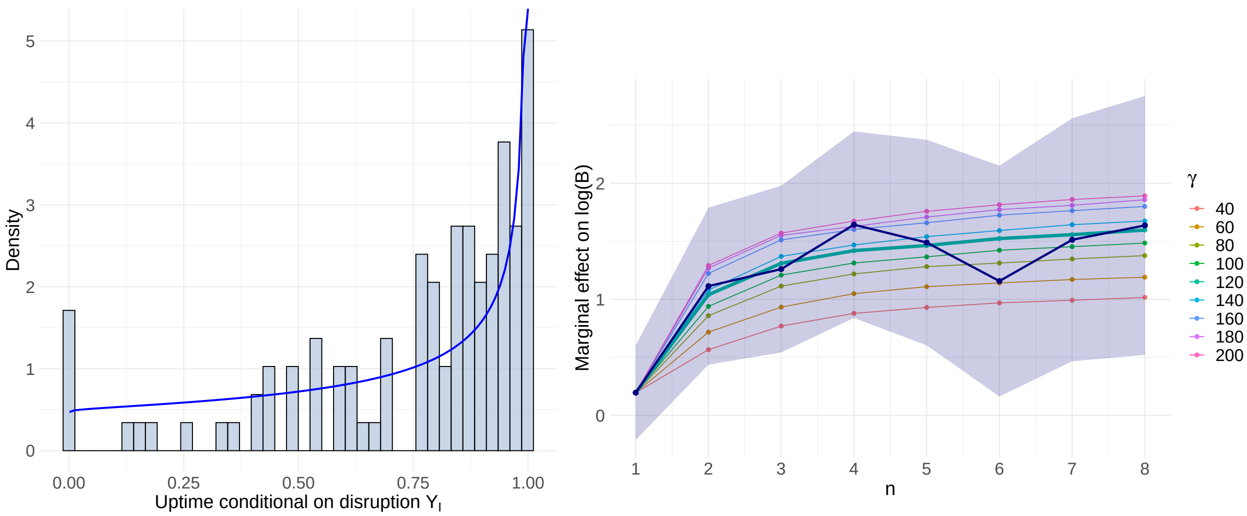

(a) Histogram of Y_i and Estimated Density

(b) Effect of n on $\log(B)$: $\frac{1}{\alpha}\log\left(\widetilde{h}(n,\gamma)\right)$

Figure A2: Panel (a) shows the histogram of uptime Y_i and the density of a Beta distribution with shape parameters 1.014 and 0.516 estimated by maximum likelihood. Panel (b) plots $\frac{1}{\alpha}\log\left(\widetilde{h}(n,\gamma)\right)$ for different values of γ, with $\frac{1}{\alpha}$ set to the estimated price elasticity ($\widehat{\alpha}_p = -1.36$) and p_d set to 0.10. The solid navy line corresponds to the coefficient estimates $\widehat{\alpha}_n$ from the aggregate demand model (Equation (19)), the shaded area corresponds to the 95% CI.

We perform a cross-validation exercise to assess the predictive performance of this model by comparing model predictions of purchased bandwidth to the data. While we observe used bandwidth B_{mt} at the market level over the timeframe 2005–2021, data on purchased bandwidth is more limited and only available at the region-pair level and over the timeframe 2010–2021. We compare purchased bandwidth Q_{rt} observed for region-pair r in period t to the purchased bandwidth predicted by our model. Given used bandwidth B_{mt}, we estimate purchased bandwidth for region-pair r as

$$\widehat{Q}_{rt} = \sum_{m\in r}\widehat{Q}_{mt} = \sum_{m\in r} B_{mt}/\widetilde{f}(n_{mt},\gamma)$$

where n_{mt} is the number of cables operating in market m in period t. We compute bounds on $\widehat{Q}_{rt}$ using two extreme values for γ (40 and 200). The results are shown in Figure A3 which plots Q_{rt} against $\widehat{Q}_{rt}$ (with γ equal to 40 and 200), both in logarithm.

The average ratio of used to purchased bandwidth (B_{rt}/Q_{rt}) in the data is 0.73 and the median is 0.77. The model predictions for the ratio closely match the data. This cross-validation exercise provides evidence that the consumer optimization problem and model parameterization yield reasonable predictions for buyers' bandwidth purchasing and usage behavior.

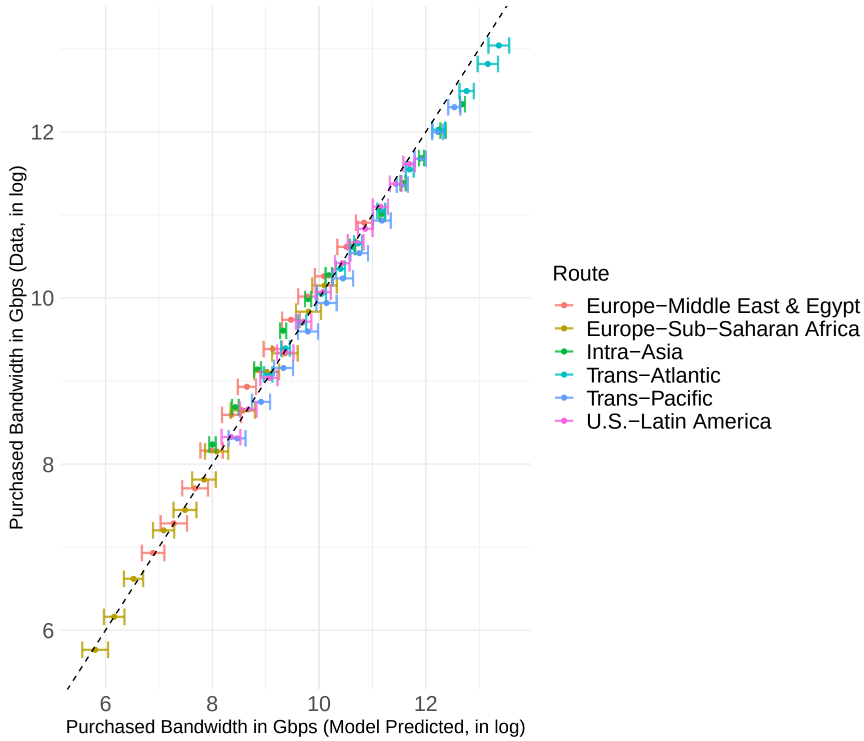

Figure A3: This figure plots purchased bandwidth Q_{rt} observed in the data for six different region pairs over 2010–2021, against purchased bandwidth predicted by our model as $\sum_{m\in r} B_{mt}/\widetilde{f}(n_{mt},\gamma)$ for γ equal to 40 and 200. The left side of each horizontal segment corresponds to γ equal 40, the right side corresponds to γ equal 200.

D Solution Method and Estimation Approaches

Solution Method. The dynamic game is solved by backward induction starting from the last period in our sample, i.e., $t = T$, market by market. In market m, denote a given element of the state space as

$$\mathcal{M}_{m,t} = (n_{mt}, d_{mt}, h_{mt}),$$

where n_{mt} is the number of cables in operation, d_{mt} is the exogenous demand state, and h_{mt} is the exogenous cost state. The latter two variables are discretized for tractability and we provide details about the discretization procedure in Sections 5.3 and 5.4. To account for market-level unobservables and the non-stationary nature of the industry, the vector of CCPs and value functions are indexed by the market and the period (in addition to the state).

We iterate over the following steps.

1. Initialize the vectors of CCPs for each market m, state $\mathcal{M}_{m,t}$, and period t, denoted $\mathbf{P}_{m,t}$. The vector $\mathbf{P}_{m,t}$ is indexed by the state $\mathcal{M}_{m,t}$ and gives the CCP of entry into the market m in state $\mathcal{M}_{m,t}$ in period t.

2. For each market m and period t, form the transition matrix from state $\mathcal{M}_{m,t}$ to state $\mathcal{M}_{m,t+1}$ *conditional* on the action played a. This transition matrix is also indexed by time because CCPs and the transition processes of exogenous states (d_{mt}, h_{mt}) are indexed by time. Two transition matrices are necessary: for incumbents conditional on staying in the market, and for potential entrants conditional on entering. Denote the transition matrices for incumbents and entrants $\mathbf{F}^i_{m,t}$ and $\mathbf{F}^e_{m,t}$ respectively. If a firm plays a terminal action (that is, a potential entrant stays out) the continuation value is zero, therefore, knowledge of this transition matrix is not necessary.

3. For each market m and period t, solve for the ex-ante value function of an incumbent. For period $t \geq T$, the ex-ante value function solves the system of equations

$$\begin{aligned} \mathbf{V}^i_{m,T} &= \boldsymbol{\pi}_{m,T} - FC_m + \beta \mathbf{F}^i_{m,T} \mathbf{V}^i_{m,T} \\ &= \left(I - \beta \mathbf{F}^i_{m,T}\right)^{-1} \left[\boldsymbol{\pi}_{m,T} - FC_m\right], \end{aligned} \tag{28}$$

where I is the identity matrix, $\boldsymbol{\pi}_{m,T}$ is a vector of variable profits in each state and FC_m is a market-specific fixed cost.[41] For period $t \leq T-1$, the ex-ante value function is obtained recursively as

$$\mathbf{V}^i_{m,t} = \boldsymbol{\pi}_{m,t} - FC_m + \beta \mathbf{F}^i_{m,t} \mathbf{V}^i_{m,t+1}\,. \tag{29}$$

4. Update the conditional choice-specific value function. Let $\mathbf{v}^e_{m,t}$ denote a vector collecting the choice-specific value function from entering in market m in period t. This vector satisfies the equality

$$\mathbf{v}^e_{m,t} = -EC_m + \beta \mathbf{F}^e_{m,t} \mathbf{V}^i_{m,t+1}\,. \tag{30}$$

5. Update the vectors of CCPs as

$$\mathbf{P}^{'}_{m,t} = \frac{\exp\left(\frac{\mathbf{v}^e_{m,t}}{\theta^\epsilon_m}\right)}{1 + \exp\left(\frac{\mathbf{v}^e_{m,t}}{\theta^\epsilon_m}\right)}\,. \tag{31}$$

If the maximum absolute difference between $\mathbf{P}_{m,t}$ and $\mathbf{P}^{'}_{m,t}$, across periods $t = 1, \ldots, T$, is less than the pre-defined tolerance level (10^{-4}), the procedure stops and $\left(\mathbf{P}^{'}_{m,t}\right)_{t=1,\ldots,T}$ is saved. If not, define updated CCPs as a convex combination of old and new CCPs $\eta \mathbf{P}_{m,t} + (1-\eta)\mathbf{P}^{'}_{m,t}$ for each player i and return to Step 2.

Because markets are independent, we can solve the model for each market separately. For our counterfactual analysis, we initialize this algorithm at a large number of starting values and iterate to a fixed point.

Estimation Approach. Given a vector of CCPs and structural parameters in iteration k,

[41] We do not model exit decisions, therefore, we rule out firm-specific private information shocks for incumbents.

denoted $(\mathbf{P}^{(k)}, \boldsymbol{\theta}^{(k)})$, we iterate over the following steps of the NPL algorithm.

1. Update the transition matrix from state $\mathcal{M}_{m,t}$ to state $\mathcal{M}_{m,t+1}$ *conditional* on the action played a. As in the solution method, we store two transition matrices: for incumbents conditional on staying in the market, and for potential entrants conditional on entering. Denote the transition matrices for incumbents and entrants $\mathbf{F}^{i,(k+1)}_{m,t}$ and $\mathbf{F}^{e,(k+1)}_{m,t}$ respectively.

2. Solve for the ex-ante value function of an incumbent *gross* of the fixed cost, denoted $\widetilde{\mathbf{V}}^{i,(k+1)}_{m,t}$.[42]

$$\begin{aligned}\widetilde{\mathbf{V}}^{i,(k+1)}_{m,T} &= \boldsymbol{\pi}_{m,T} + \beta \mathbf{F}^{i}_{m,T} \widetilde{\mathbf{V}}^{i,(k+1)}_{m,T} \\ &= \left(I - \beta \mathbf{F}^{i,(k+1)}_{m,T}\right)^{-1} \boldsymbol{\pi}_{m,T} \,,\end{aligned} \tag{32}$$

and

$$\widetilde{\mathbf{V}}^{i,(k+1)}_{m,t} = \boldsymbol{\pi}_{m,t} + \beta \mathbf{F}^{i}_{m,t} \widetilde{\mathbf{V}}^{i,(k+1)}_{m,t+1} \,. \tag{33}$$

3. Estimate the structural parameters $\boldsymbol{\theta}^{(k+1)}$ using a pseudo-likelihood estimator where the probabilities that a potential entrant enters in market m and period t is given by

$$\frac{\exp\left(\left\{-EC_m + \beta \mathbf{F}^{e,(k+1)}_{m,t} \widetilde{\mathbf{V}}^{i,(k+1)}_{m,t+1} - \frac{\beta}{1-\beta} FC_m\right\} \Big/ \theta^{\epsilon}_m\right)}{1 + \exp\left(\left\{-EC_m + \beta \mathbf{F}^{e,(k+1)}_{m,t} \widetilde{\mathbf{V}}^{i,(k+1)}_{m,t+1} - \frac{\beta}{1-\beta} FC_m\right\} \Big/ \theta^{\epsilon}_m\right)} \,. \tag{34}$$

4. If the maximum absolute difference between $\boldsymbol{\theta}^{(k)}$ and $\boldsymbol{\theta}^{(k+1)}$ and between $\mathbf{P}^{(k)}_{m,t}$ and $\mathbf{P}'_{m,t}$ (based on Equation (34) under $\boldsymbol{\theta}^{(k+1)}$) is less than the tolerance level, stop the procedure. If not, return to step 1 using $\mathbf{P}^{(k+1)}_{m,t} = \mathbf{P}'_{m,t}$ and $\boldsymbol{\theta}^{(k+1)}$.

In some cases, the NPL algorithm may fail to converge if the fixed point corresponding to the data generating process is unstable (Pesendorfer and Schmidt-Dengler (2010)). To address this concern, we implement several alternative estimators: two-step estimators (1-PML and 1-MD) and the spectral algorithm recently proposed by Aguirregabiria and Marcoux (2021). The latter estimator does not iterate on the best-response mapping to attain a fixed-point, rather, it solves for the root of a nonlinear system of equations by a quasi-Newton method. As a consequence, the spectral approach can find unstable fixed points that would be unattainable by the NPL algorithm.

Define $\phi(\mathbf{P}) \equiv \mathbf{P} - \Psi(\mathbf{P}, \hat{\theta}(\mathbf{P}))$, where $\hat{\theta}(\mathbf{P})$ is the pseudo-likelihood maximizer given input CCP vector $\mathbf{P}$ and $\Psi(.,.)$ is the best-response mapping as a function of input CCP and structural parameters. To find the solution(s) to $\phi(\mathbf{P}) = 0$, spectral approaches are

[42] We compute the value function gross of the fixed cost in order to be able to express the probability of entry as a function of the fixed cost and maximize the pseudo-likelihood. If we were to use Equations (28) and (29), then FC_m would be subsumed in the value function and would not appear as a argument in the pseudo-likelihood.

particularly useful because they do not require computing the large-dimensional Jacobian $\nabla\phi(\mathbf{P})$ as would be required in Newton's method.[43] In practice, we replace the updating step above (step 4, where $\mathbf{P}^{(k+1)} = \Psi(\mathbf{P}^{(k)}, \hat{\theta}(\mathbf{P}^{(k)})))$ by

$$\mathbf{P}_{m,t}^{(k+1)} = \mathbf{P}_{m,t}^{(k)} - \zeta_k \phi(\mathbf{P}_{m,t}^{(k)})\,,$$

where ζ_k, the spectral step, equals

$$\zeta_k = \frac{\Delta\mathbf{P}^{(k)\prime}\Delta\mathbf{P}^{(k)}}{\Delta\mathbf{P}^{(k)\prime}\Delta\phi(\mathbf{P}^{(k)})}\,,$$

and $\Delta\mathbf{P}^{(k)} = \mathbf{P}^{(k)} - \mathbf{P}^{(k-1)}$, $\Delta\phi(\mathbf{P}^{(k)}) = \phi(\mathbf{P}^{(k)}) - \phi(\mathbf{P}^{(k-1)})$. We set the initial value ζ_0 to one.[44]

E Addressing the Endogeneity of Entry

If firms strategically time entry decisions to coincide with demand shocks that are unobserved by the econometrician, the estimated effects of the number of cables in Equation (19) may be biased and may not reflect buyers' preference for diversification. We address this endogeneity issue using two distinct approaches.

Instrumental Variables. We use the occurrence of cable disruptions as an exogenous shifter of entry and subsequent market structure: we find support in the data that markets hit by cable disruptions experience significantly less entry in the next period, two likely reasons being that: (1) repairing disruptions requires the same specialized cable ships that are also used in the laying of new cables; and (2) some cable disruptions are caused by tectonic and weather events that also hinder the laying of cables. Industry reports such as Swinhoe (2022b,a) and Dzieza (2024) document the limited and aging fleet of cable ships; as of 2022 there were roughly 60 cable ships in the world, with no newly-built cable ships launching between 2004 and 2010 and only five such ships launching between 2011 and 2020.[45] In the last few years, cable ships laying new cables were booked at least 24 months in advance and industry experts expect that the industry may move toward long-term exclusive contracts (akin to deep water oil rigs).[46]

Table A3 present the results of this analysis. In columns (1) to (6), we focus on the time-

[43] Newton's method updates the CCP as $\mathbf{P}_{m,t}^{(k+1)} = \mathbf{P}_{m,t}^{(k)} - [\nabla\phi(\mathbf{P}_{m,t})]^{-1}\phi(\mathbf{P}_{m,t}^{(k)})$.

[44] The spectral approach replaces the inverse of the Jacobian matrix $[\nabla\phi(\mathbf{P})]^{-1}$ by ζ_k. As shown in Aguirregabiria and Marcoux (2021), $1/\zeta_k$ approximates a Rayleigh quotient of $\nabla\phi(\mathbf{P})$: it is a weighted average of the eigenvalues of this matrix.

[45] For example, when Vietnam experienced five cable disruptions by routine causes in the early part of 2024, repairs lasted 6 months due to constraints in ship availability.

[46] An example of how the limited capacity of the cable ship fleet binds is the Maple Leaf Fiber cable. Conceived in 2018 as a cable with an underwater segment connecting Toronto and Kingston (under lake Ontario) and a terrestrial segment between Kingston and Montreal, the project was amended in 2022 to use terrestrial cables for both segments (Swinhoe, 2022b).

frame 2013–2021 and drop earlier years (2005-2013) for which data on cable disruptions is not available. In all specifications, we control for the logarithm of one plus the number of cables operating in the market. Column (1) shows results based on an OLS regression. Columns (2) and (3) show the first and second stages of an IV regression (model IV-1) where only price endogeneity is addressed. Columns (4) to (6) show the first stages and second stage of an IV regression (model IV-2) where both price and entry endogeneity are addressed. Cable entry in period t is instrumented by lagged cable faults (in $t-1$). The first stage regression for the number of cables (Column (5)) shows that the number of cable faults in $t-1$ has a negative and statistically significant effect on the number of cables operating in current period t, indicative of delays in entry. The estimated effect of the number of cables (column (6)) on demand equals 0.920 and is similar to the baseline estimate (0.977 in column (3)). This result alleviates concerns about remaining unobserved demand shocks correlated with entry decisions.

Difference-in-difference. Alternatively, we address endogeneity concerns around entry decisions using a difference-in-difference like strategy. In Column (7) (model IV-3), we conduct the same IV regression as in columns (2) and (3), but at an *annual* level from 2005 to 2021.[47] We control for region-year and market fixed effects, assuming common demand growth rates across markets within the same region pair (e.g., Trans-Atlantic, Trans-Pacific, Europe-Asia), and identify the effect of entry on demand from growth rate differences between markets in the same region pair. Despite the smaller sample size (one-fourth of the full quarterly sample), the estimated price elasticity and the effect of the number of cables are comparable to other specifications. We also experimented with specifications that include the number of cables as categorical variables as in Equation (19), and found consistent results.

F Extensions

F.1 Time-to-build

In this section, we extend the industry model to account for the fact that cables typically take more than one period (3 months) to build. A complete specification includes two features: a function mapping cable and market characteristics (e.g., geography, cable length) to time-to-build and a modification of the state space to track cables under construction and their expected completion time. While the latter is conceptually simple, it significantly increases the state space size, complicating empirical implementation. For example, with a two-year time-to-build, the entry probability for a cable in a quarter depends on how many cables began construction in the past seven quarters, adding seven dimensions. This full

[47] We omit the first-stage regression for price in model IV-3, as the estimated effects are similar to those reported in Column (2).

Table A3: Robustness Checks: Used Bandwidth Demand

	OLS	IV-1		IV-2			IV-3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		1st-Stage: Price	2nd-Stage	1st-Stage: Price	1st-Stage: Cables	2nd-Stage	2nd-Stage
Bandwidth Price (10G, log)	-0.129		-1.309			-1.207	-1.373
	(0.0583)		(0.310)			(0.289)	(0.725)
Number of undersea cables (log)	1.293	-0.259	0.977			0.920	0.854
	(0.307)	(0.170)	(0.359)			(0.466)	(0.318)
Demand factors							
Fixed Broadband Subscriptions (log)	0.251	0.126	0.413	0.142	-0.0581	0.395	0.417
	(0.113)	(0.0704)	(0.140)	(0.0687)	(0.0283)	(0.135)	(0.147)
GDP (log)	1.241	0.206	1.497	0.174	0.0898	1.480	1.336
	(0.397)	(0.181)	(0.408)	(0.172)	(0.0552)	(0.404)	(0.300)
Aggregate trade flow (log)	-0.0179	0.0201	0.00541	0.0203	-0.00685	0.00234	0.00333
	(0.0439)	(0.0279)	(0.0661)	(0.0286)	(0.00437)	(0.0650)	(0.0781)
Number of data/cloud centers (log)	0.123	-0.00400	0.129	0.000127	-0.0211	0.126	0.225
	(0.0903)	(0.0829)	(0.142)	(0.0825)	(0.00658)	(0.137)	(0.134)
Electricity price (log)		0.0780		0.0783	-0.0101		
		(0.0250)		(0.0251)	(0.00538)		
% change in Electricity price		-0.248		-0.249	0.0000245		
		(0.0376)		(0.0377)	(0.00735)		
Cable faults in $t-1$ (log)				-0.0332	-0.00812		
				(0.0182)	(0.00336)		
Country Pair FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region Pair × Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Weak Identification test			24.03			12.20	4.194
Endogeneity test (p-value)			32.2 (0)			29.4 (0)	4.10 (0.043)
R^2	0.98		0.96			0.97	0.96
Adjusted R^2	0.98		0.96			0.96	0.94
Observations	2673	2676	2676	2676	2676	2676	1016

Note: In columns (1) to (6), the unit of observation is a country pair by quarter. In columns (7), the unit of observation is a country pair by year. Standard errors are clustered at the country pair level. Distance corresponds to the bilateral distances between countries, calculated as a weighted arithmetic average of the geodesic distances between the main cities in these countries, where population weights are used. For unilateral variables (GDP, fixed broadband subscriptions, data centers, electricity prices), we restrict the coefficient to be the same for both countries in a pair.

specification is beyond the scope of this paper.

Instead, we use a simplified characterization that captures the core economics without overwhelming estimation. Each new cable entrant joins the pool of cables under construction, transitioning to market competition at a constant, exogenous, market-specific rate, ζ_m. Thus, potential entrants only need to track the number of cables under construction, e_{mt}. The expected number of competitors completing construction next period is $\zeta_m e_{mt}$ (from a binomial distribution $\mathcal{B}(e_{mt}, \zeta_m)$), and a new entrant expects to wait $1/\zeta_m$ periods to complete construction and begin operations.

In this setting, the only part of the industry model that needs to be modified are the dynamics. The bandwidth demand model, the static competition game, and per-period profits are unchanged, as are the transitions for the demand and cost states d_{mt} and h_{mt}.

First, the industry state becomes

$$\mathcal{M}_{mt} = (n_{mt}, e_{mt}, d_{mt}, h_{mt}), \tag{35}$$

where the additional component e_{mt} represents the number of under-construction cables in market m in period t. We set ζ_m such that the expected wait time under construction is 2 years

Second, while the profit Equation (9) is unchanged, we note that entry costs EC_{mt} are paid in the period that the cable enters construction, and variable profits $VP(\mathcal{M}_t)$ are zero until construction is completed.

Third, the choice-specific value function for the potential entrant then becomes

$$v_{i,t}^{\alpha}(a_{it}, \mathcal{M}_t) = \begin{cases} -EC_{mt} + \beta\, \mathbb{E}_t \left[V_{i,t+1}^{\alpha,e} \left(\mathcal{M}_{t+1} \right) \right] & \text{if} \quad a_{it} = 1 \\ 0 & \text{if} \quad a_{it} = 0 \end{cases}, \tag{36}$$

where the (next-period) value function for an under-construction cable $V_{i,t+1}^{\alpha,e}$ is

$$V_{i,t+1}^{\alpha,e}\left(\mathcal{M}_{t+1}\right) = \zeta_m \beta\, \mathbb{E}_{t+1} \left[V_{i,t+2}^{\alpha} \left(\mathcal{M}_{t+2} \right) \right] + (1-\zeta_m)\beta\, \mathbb{E}_{t+1} \left[V_{i,t+2}^{\alpha,e} \left(\mathcal{M}_{t+2} \right) \right], \tag{37}$$

and $V_{i,t+2}^{\alpha}$ is the value function of an incumbent cable operator. The estimates of δ_ϵ remain qualitatively similar to the baseline, although their magnitude is smaller reflecting the fact that the continuation value from entering (Equation (36)) is lower with time-to-build.

F.2 Exogenous Cable Retirement

As in the previous section, we can introduce an alternative extension to the dynamic game: the retirement of cables after a finite time horizon. This addition to the model brings it closer to what is likely to be the long-run equilibrium in the industry (though as discussed in Section 3, not the years covered in our data). Industry consensus is that cables have roughly a 25 year lifespan (this is the length of a typical FCC license for a cable landing in the US), though some cables have been in service for longer and have received regulatory extensions.

As in Section F.1, fully specifying cable retirement leads to a curse of dimensionality. For example, assuming cables retire after 25 years, the value of entering a market depends on the age distribution of incumbent cables. A market with 1-, 5-, and 7-year-old incumbents differs from one with 22-, 23-, and 24-year-old incumbents. Thus, the value function would depend on both the number and age distribution of incumbents.

To maintain tractability, we adopt a simplified retirement model. Each cable faces an exogenous and constant retirement probability, χ, per period.[48] This implies an expected lifespan of $1/\chi$ periods. Retirement is a terminal state, with a retired cable replaced by a new potential entrant.

Under this assumption, the model described in Section 4 can be used with a few minor modifications. The dynamic state remains

$$\mathcal{M}_{mt} = (n_{mt}, d_{mt}, h_{mt}), \tag{38}$$

and the introduction of the exogenous exit rate χ alters the expectation of future competition, as (without entry) $\mathbb{E}_t[n_{t+1}] = (1-\chi)n_t$.[49] The choice-specific value function for incumbent firms becomes:

$$v^{\alpha}_{i,t}(a_{it}, \mathcal{M}_t) = \begin{cases} \pi_i(1, \mathcal{M}_t) + \beta(1-\chi)\,\mathbb{E}_t\left[V^{\alpha}_{i,t+1}\left(\mathcal{M}_{t+1}, \boldsymbol{\epsilon}_{i,t+1}\right)\right] + \beta\chi EV_{mt} & \text{if} \quad a_{it} = 1 \\ \pi_i(0, \mathcal{M}_t) & \text{if} \quad a_{it} = 0 \end{cases}, \tag{39}$$

where it can be seen that even if an incumbent chooses to remain active ($a_{it} = 1$), they will exogenously exit and receive EV_{mt} (which is normalized to zero) with probability χ. We set χ so that the expected lifespan of a cable is 25 years. This specification does not significantly change the estimates of δ_ϵ relative to the baseline (e.g., the 1-PML estimate $\widehat{\delta_\epsilon}$ equals 0.245 under $\delta_{FC} = 0.2\%$).

F.3 Asymmetries in Marginal Costs

In this section, we allow for heterogeneity in marginal costs across cables with differing vintages (i.e. built at different times). We consider a setup with J different vintages. Suppressing the market and time subscripts for legibility, the number of incumbent cables belonging to vintage j is denoted n_j, with $n = \sum_J n_j$. The marginal costs of bandwidth for each type are likewise denoted c_j; we consider a type-symmetric equilibrium of the static Cournot game and denote the (per-firm) quantity produced by firms of type j by q_j, with the market-period identity $Q = \sum_j n_j q_j$.

[48] While the exit probability could vary by market (χ_m), limited exit events in our data prevent empirical estimation of market-specific probabilities.

[49] The state space is larger under this extension because n_t can potentially decrease, whereas under the baseline model without exit, n_t is weakly increasing over time.

A firm of type j chooses the optimal amount of bandwidth to supply by satisfying the first-order condition (omitting the market and time subscripts)

$$p - c_j + q_j \frac{\partial p}{\partial q_j} = 0\,,$$

which can be rewritten as

$$\frac{p - c_j}{p} = -\frac{q_j}{Q}\frac{Q}{p}\frac{\partial p}{\partial q_j} = -\frac{q_j}{Q}\frac{1}{\epsilon}\,,$$

where ε is the price elasticity of demand. Summing over all firms yields

$$\sum_{j=1}^{J} n_j \left(\frac{p - c_j}{p}\right) = -\frac{1}{\epsilon}\,, \tag{40}$$

which can be rearranged to express equilibrium prices as a function of the number of firms and marginal costs by type:

$$p = \frac{1}{n + \varepsilon^{-1}} \sum_{j=1}^{J} n_j c_j\,. \tag{41}$$

Noting that p, n_j, and ε are either observed in the data or estimated objects, one can then recover the marginal costs by firm type $\{c_j\}_{j \in J}$ by exploiting the cross-sectional and time-series variation in the vintage distribution of incumbent cables.

In practice, we set J to 2 and assign cables to vintages with the following cutoff

$$j(i) = \begin{cases} 1 & \text{if} \quad t_i \leq \widetilde{t} \\ 2 & \text{if} \quad t_i > \widetilde{t} \end{cases}\,,$$

where t_i is the period cable i entered service, and $\widetilde{t}$ is a cut-off period. We assume that marginal costs are the sum of market-specific δ_m, period-specific δ_t, and vintage-specific fixed effects $\widetilde{\delta}_j$, for $j \in \{1, 2\}$. This allows markets with different physical characteristics (e.g. distance, propensity for cable faults) to have different marginal costs while restricting the within-market cost difference between different vintages of cables to be common across all markets. Kalouptsidi (2014) uses a similar approach for ship ages. We estimate the following equation via ordinary least squares (denoting the error term by ν_{mt}):

$$p_{mt}\left(\frac{n_{mt} + \epsilon^{-1}}{n_{mt}}\right) = \delta_m + \delta_t + \frac{n_{1,mt}}{n_{mt}}\widetilde{\delta}_1 + \frac{n_{2,mt}}{n_{mt}}\widetilde{\delta}_2 + \nu_{mt}\,. \tag{42}$$

Table A4 shows the estimation results for different choices of the cutoff quarter $\widetilde{t}$ (e.g., 2010 corresponds to 2009-Q4). We do not find evidence of significant heterogeneity in costs across cables of different vintages, once we control for market and period fixed effects. This is related to the discussion in Section 5.4 suggesting that many technological upgrades (that affect the marginal cost of using current bandwidth or lighting additional bandwidth) can be installed on equipment at the landing stations and, therefore, benefit incumbent cables as well as new entrants.[50]

Table A4: Heterogeneity in Marginal Costs

	Dependent variable: $p_{mt}\left(\frac{n_{mt}+\epsilon^{-1}}{n_{mt}}\right)$		
	$\widetilde{t}=2005$	$\widetilde{t}=2010$	$\widetilde{t}=2015$
$n_{1,mt}/n_{mt}$	13378.0	13443.2	12761.0
	(2403.8)	(2324.4)	(2235.4)
$n_{2,mt}/n_{mt}$	12849.2	12600.8	14379.9
	(3495.7)	(3559.8)	(3879.8)
Market FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
R^2	0.776	0.776	0.776
Adjusted R^2	0.763	0.763	0.763
Within R^2	0.009	0.009	0.009
Observations	4,777	4,777	4,777

Note: The unit of observation is the market (country pair) by quarter. Standard errors (in parenthesis) are clustered at the market level. The depend variable is expressed in $US.

F.4 Time-varying Investment Costs

To more accurately capture technological changes affecting investment costs, we allow entry costs (EC_{mt}) to vary over time. Using cable construction cost data (available for 47% of cables), we predict entry costs for all markets, adding a time trend to account for decreasing costs. Including the time trend slightly improves the fit of the construction cost regression (R^2 increases from 0.76 to 0.83), and entry costs decline modestly by about 1.8% annually over the sample period.

[50] Fiber-optic transmission system upgrades are technologies and methods used to increase the capacity and improve the performance of *existing* fiber-optic transmission systems. Examples include Wavelength division multiplexing, a technique that allows multiple wavelengths to be transmitted over a single fiber-optic cable at the same time; Optical amplifiers, devices that boost the strength of the light signals that are transmitted over a fiber-optic cable; Polarization-multiplexing, a technique used to transmit two different data streams over a single fiber using different orientations of light waves (or polarization); Space-division multiplexing, a technology used to increase the capacity of the cable by using multiple cores within a single optical fiber. These techniques are generally considered to be relatively low-cost as they do not involve significant capital investments. Software-defined networking can be used to manage the traffic on the cable by providing a centralized control plane that can make decisions about how to route traffic based on real-time information about the network's state. This allows for greater flexibility in managing the network and can enable more efficient use of network resources.

The model is estimated with time-varying entry costs EC_{mt}, while maintaining our assumption that fixed costs are proportional to the level of entry costs $FC_{mt} = \delta_{FC} EC_{mt}$. We find that our key estimates for the dynamic parameters remain quantitatively similar to our original findings: e.g., the 1-PML estimate $\widehat{\delta}_\epsilon$ equals 0.268 under $\delta_{FC} = 0.2\%$. The results of the counterfactual exercise under these more flexible entry costs align with our baseline predictions. Notably, supplier diversification drives an important share of investment in new cables. Because entry costs are decreasing over time, the welfare cost from excess entry is reduced, leading to a smaller relative magnitude of business-stealing to diversity distortions. Overall, these results suggest that the core economic mechanisms and insights of our analysis are robust to allowing for more flexible cost dynamics.

G Determinants of Marginal Costs

We investigate how marginal costs vary with cable characteristics and other cost shifters. The following specification for the marginal cost function is used

$$\log(mc_{mt}) = \gamma_0 \mathbf{W}_{mt} + \gamma_q \log(q_{mt}) + \eta_m + \eta_{r(m)t} + \omega_{mt}\,, \tag{43}$$

where $\mathbf{W}_{mt}$ are exogenous cost shifters, q_{mt} is the quantity supplied by a given firm (or "purchased" bandwidth), η_m and $\eta_{r(m)t}$ are market fixed effects and region by time fixed effects, and ω_{mt} are unobserved cost shocks at the market-period level. Purchased bandwidth per firm q_{mt} is equal to Q_{mt}/n_{mt}, where Q_{mt} is the total purchased bandwidth derived from the observed used bandwidth B_{mt} and the used-to-purchased bandwidth ratio $\widetilde{f}(n_{mt}, \gamma)$ with γ set to 120 (see Appendix C).

This specification allows marginal costs to depend on quantity to capture potential economies of scale. We estimate this specification by OLS and show the results in Table A5. As expected, we find a positive correlation between marginal costs and electricity prices, cable length, and the number of (contemporaneous) cable faults. In specifications (5) and (6), marginal costs are also decreasing in the amount of bandwidth supplied by a given cable, suggesting the presence of economies of scale. However, this effect is not robust to the inclusion of market-level fixed effects (specification (7)).[51]

H Model Fit and Assumptions

In this section, we discuss our modeling assumptions concerning the process of demand and cost states, the conduct model, and the independence of markets. We also examine the plausibility of our cost estimates and the fit of the model by comparing predictions of the model to the data.

[51] Because q_{mt} may be correlated with the unobserved cost shock, we also run specification (7) instrumenting q_{mt} with demand-side shifters. The results remain quantitatively similar.

Table A5: Determinants of Marginal Costs

	Dependent variable: Marginal cost						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Purchased bandwidth	-0.386				-0.373	-0.137	-0.007
	(0.008)				(0.007)	(0.005)	(0.015)
Electricity Price		0.535			0.255	0.086	0.194
		(0.025)			(0.021)	(0.030)	(0.036)
% Change in Electricity Price		-0.234			-0.202	-0.159	-0.706
		(0.072)			(0.058)	(0.100)	(0.095)
Cable length			0.588		0.556	-0.035	
			(0.017)		(0.015)	(0.017)	
Number of cable faults				0.127	-0.039	-0.002	0.009
				(0.034)	(0.024)	(0.014)	(0.013)
Region Pair FE	No	No	No	No	No	Yes	No
Time FE	No	No	No	No	No	Yes	No
Region Pair $\times$ Time FE	No	No	No	No	No	No	Yes
Market FE	No	No	No	No	No	No	Yes
R^2	0.348	0.166	0.196	0.004	0.606	0.850	0.976
Adjusted R^2	0.348	0.165	0.196	0.003	0.606	0.847	0.969
Within R^2						0.112	0.029
Observations	4,777	4,493	4,777	3,714	3,430	3,430	3,430

Note: The unit of observation is the market (country pair) by quarter. All variables are in log (except change in electricity prices and number of cable faults). The number of cable faults is only available starting in 2013. Purchased bandwidth corresponds to the amount supplied by a single firm. Cable length is averaged over all cables operating in a given market.

Transitions of Aggregate States. First, we consider the estimated transition processes of the demand and cost states (d_{mt}, h_{mt}). These constitute important inputs into the dynamic game. We use the fitted AR(1) processes to simulate these two state variables forward and compare simulation results to the realization of the states in the data. Figure A4 shows the results for the market "United Kingdom-United States." The left panel shows the demand state d_{mt}, and the right panel shows the cost state h_{mt} in the data, along with 95% confidence intervals from our simulations (dotted lines) and the median simulated value (dashed line).[52] We show similar simulation exercises for a subsample of markets in Figures A9 and A10 of Appendix I.

Overall, the estimated transition processes capture the time series variation and the heterogeneity across markets well. For the demand state d_{mt}, we find a positive trend for some markets (e.g., $\exp(d_{mt})$ increased seven-fold over the sample period for the market France–UK), while in other markets, decreasing prices (P_{mt}) and, to a lesser extent, increasing diversity (n_{mt}) explain most of the rise in B_{mt}, leading to more muted trend in the exogenous

[52] The last period in the data is 2021-Q4. For the simulations, we iterated forward until 2025.

Figure A4: Demand State and Cost State Over Time for the Market United Kingdom-United States

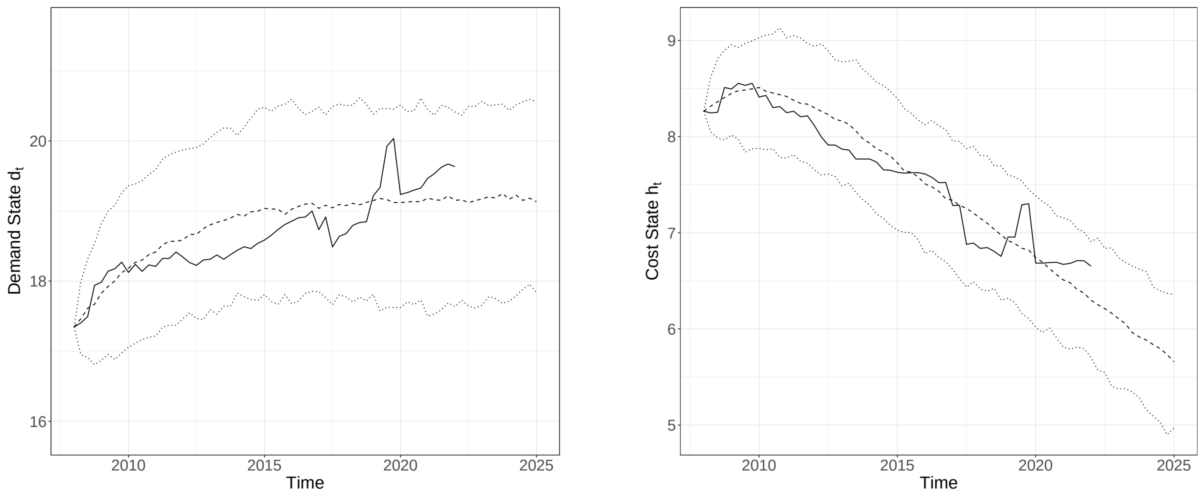

(a) d_{mt} (b) h_{mt}

Notes: The figures show the estimated demand d_{mt} and cost states h_{mt} over time (solid lines). To demonstrate the fit of our estimated transition processes, the dotted lines give the 95% confidence interval from simulations using the transition processes. The dashed line shows the median simulated value.

part of demand (e.g., Japan–US).[53]

Magnitude of Cost Estimates. Second, we compare the magnitude of our cost estimates to industry reports and cable financial investment plans (Reverdy and Skenderoski (2015), Seixas (2015)). These reports divide costs into capital expenditures (CAPEX), payable to manufacturers and installation suppliers, and operational expenditures (OPEX), incurred annually over the cable's lifespan. OPEX to CAPEX ratios in these reports range from 4% to 7% depending on system size and characteristics. For comparison, we calculate the ratio of annual costs (variable and fixed) to entry costs. Assuming $\delta_{FC} = 0.2\%$, the average ratio is 8%, with a median of 3%. For $\delta_{FC} = 0.8\%$, the mean ratio is 10%, and the median is 5%.[54] Overall, our cost estimates align with the magnitudes in industry reports.

Model Fit. Third, given our structural parameter estimates, we solve for the model equilibrium and compare the model predictions to the actual data. This is a useful way of evaluating the assumptions embedded in the model, in particular, those concerning information, timing,

[53]Another potential explanation for the stability of d_{mt} in some markets is the growth in content delivery networks (CDNs) over our sample period. CDNs are networks of geographically distributed data centers designed to enhance the delivery of web content by strategically placing copies of data in multiple server locations around the world. In this sense, CDNs may reduce the reliance on data transport via undersea cables by substituting local storage for transport. In our demand specification, the growth in CDNs is captured by the region-time fixed effects $\gamma_{r(m)t}$.

[54]Outlier markets exist where this ratio exceeds one. For instance, geographically close markets like Denmark-Sweden or Indonesia-Singapore have low entry costs due to short cables but high variable costs due to heavy bandwidth traffic.

functional forms, etc. The equilibrium CCP (solved for) are used to simulate the industry forward from the initial industry state and predict the expected number of cables in the last period of the sample. Figure A5 shows the number of cables in the last period in the data against the (simulated) expected number of cables predicted by the equilibrium of the dynamic game. We find that the model performs well, predicting slightly too much entry into markets that have only one cable in the data, but otherwise matching the number and distribution of cables very closely.

Figure A5: Model Fit for the Number of Cables in the Last Period

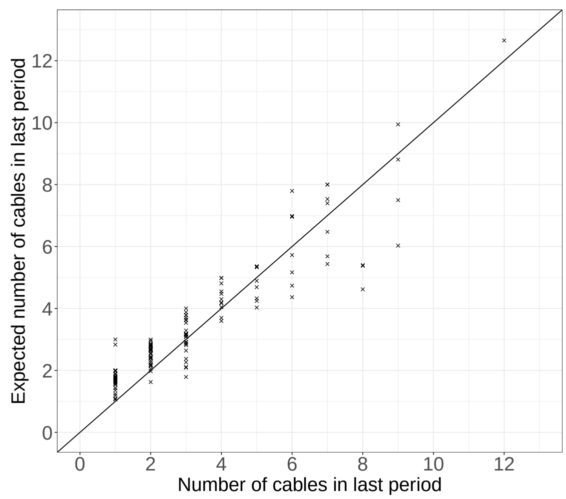

Notes: This figure provides an illustration of how the dynamic model matches the data by plotting the number of cables in the data (horizontal axis) against the simulated number of cables (vertical axis) in the last period for each market. The model fit is good, with a slight tendency to over-predict entry into those markets where only one cable is present in the data.

Testing Conduct. Fourth, we test the conduct assumption (symmetric Cournot) using the approach introduced in Pakes (2017).[55] We leverage the dual facts that one can obtain an independent estimate of the markup from the demand parameters, and that our conduct assumption implies, via the first-order condition, that this markup has a coefficient of one in the quantity-setting equation. Rewriting firms' first-order condition (Equation (22))

$$p_{mt} = mc_{mt} - \frac{\partial p_{mt}}{\partial q_{imt}} q_{imt} = \gamma_0 \mathbf{W}_{mt} + \omega_{mt} - \frac{\partial p_{mt}}{\partial q_{imt}} q_{imt} \,, \tag{44}$$

where $\mathbf{W}_{mt}$ are exogenous cost factors (as in Table A5), ω_{mt} are unobserved cost shocks, and $-\frac{\partial p_{mt}}{\partial q_{imt}} q_{imt}$ is the markup. The latter variable can be calculated directly using the demand estimates. Note that this variable is endogenous as it is a function of both unobserved

[55]Due to limitations in the data, we cannot use more involved approaches (e.g., non-nested testing, model selection): in particular, without cable-level market shares, it is difficult to obtain price elasticities under a differentiated Bertrand assumption.

Table A6: Fit of Equilibrium Quantity-Setting Equation

	Dependent variable: Price	
	Estimate	S.E.
Markup	1.358	(0.238)
Electricity Price	170.119	(69.520)
% Change in Electricity Price	-415.521	(1172.760)
Number of cable faults	167.727	(422.788)
Cable length	-0.057	(0.085)
R^2	0.756	
Adjusted R^2	0.755	
Observations	2731	

Notes: The unit of observation is the market (country pair) by quarter. Standard errors are clustered at the market level and shown in parenthesis. All variables are in level. Cable length is averaged over all cables operating in a given market.

demand and cost shocks. The test is implemented as follows: (1) use the demand estimates to construct the markup term $-\frac{\partial p_{mt}}{\partial q_{imt}} q_{imt}$, (2) regress prices on observed cost factors and the markup, where the markup is instrumented using exogenous demand shifters (in practice we use GDP, cloud/data centers, and broadband subscriptions), (3) test statistically whether the estimated coefficient for the markup term is equal to one. The regression results are shown in Table A6. The fit of the regression is high ($R^2 = 0.76$). The markup coefficient is 1.36 (se = 0.24), and we cannot reject the null hypothesis of equality to one (the $p-$value is 0.14). While this result does not constitute conclusive evidence in favor of our conduct assumption, it is reassuring that the data is not at odds with the Nash in quantities assumption.

Independent Markets. Finally, we discuss the assumption of independent markets. In the industry model, cables in each market are assumed to play a separate Markov Perfect equilibrium, unaffected by neighboring markets. One concern is that demand might be correlated across neighboring markets. This is mitigated by controlling for market-level unobserved heterogeneity and regional-time trends in demand estimation. We allow for macroeconomic shocks at the region-pair level that simultaneously shift demand for neighboring markets (e.g., U.S.-U.K., U.S.-France, U.S.-Ireland).

Another concern is that demand in one market might be partially carried via indirect paths through a third country.[56] Since we observe only aggregate bandwidth flows between country pairs and not exact paths, we cannot directly verify this assumption. Industry experts confirm that data on the exact routing of bandwidth would be extremely difficult if not

[56] This issue also arises in airline industry studies when defining product markets. Many studies filter markets to select homogeneous itineraries where most trips are non-connecting, nonstop trips.

Table A7: Test of Market Independence

	Dependent variable: Log Used Bandwidth	
	Estimate	S.E.
Lag of Log Used Bandwidth	0.958	(0.003)
Focal Market Cable Count	0.010	(0.003)
Neighboring Country Cable Count	-0.000	(0.003)
R^2	0.976	
Adjusted R^2	0.975	
Observations	7535	

Notes: The unit of observation is the market (country pair) by quarter, and the regression includes region-pair by year fixed effects. Standard errors are shown in parentheses.

impossible to compile. However, we provide suggestive evidence supporting this assumption. Specifically, we examine the correlation between bandwidth growth in a focal market and the number of cables in both the focal market and neighboring markets. We expect a positive correlation between the number of cables and bandwidth growth in the focal market. Under the independence assumption, the number of cables in neighboring markets should not correlate with bandwidth growth in the focal market, which our findings support.

To test this, we identify for each focal market the largest country to which it is connected, termed the "neighboring" country. "Largest" is defined by total bandwidth connected to the country in 2021. For instance, for U.S.-France, the neighboring country is the UK. We regress bandwidth growth in the focal market on (1) the number of cables in the focal market and (2) the number of cables connecting the focal market countries to their neighboring country. The results are reported in Table A7.[57]

Second, we compare the potential capacity (theoretical design capacity at maximum utilization, available for a subset of cables in 2021) in a focal market to the used bandwidth. On average, potential capacity is four times larger than bandwidth demand per market, consistent with aggregate figures in Table 1. This indicates direct paths are not capacity-constrained in serving demand.[58]

I Supplementary Tables and Figures

[57]This analysis uses a larger sample size than the main text, as it requires only data on bandwidth and number of cables, not bandwidth prices.

[58]Transiting through a third country raises costs due to IP transit fees (akin to tolls) paid to the provider transmitting the data. Direct paths via undersea cables, if available, are more cost-efficient.

Table A8: Descriptive Statistics on Cable Faults

Year	Cables	Faults	Propensity (%)
2013	301	7	2.3
2014	316	8	2.5
2015	325	7	2.2
2016	341	7	2.1
2017	356	24	6.7
2018	374	17	4.5
2019	397	15	3.8
2020	418	42	10.0
2021	435	27	6.2
2022	459	14	3.1

Notes: Columns 2 and 3 show the number of active cables and faults reported each year. Column 4 shows the annual propensity for a cable to suffer a fault (in percentage terms).

Table A9: Transition Process for the Demand State d_{mt}

	(1)	(2)	(3)
Demand state in $t-1$		0.928	0.928
		(0.00565)	(0.00571)
Time trend	0.0278		0.000426
	(0.00306)		(0.00117)
Market-level Intercept	Yes	Yes	Yes
R^2	0.89	0.98	0.98
Adjusted R^2	0.89	0.98	0.98
Observations	4644	4644	4644

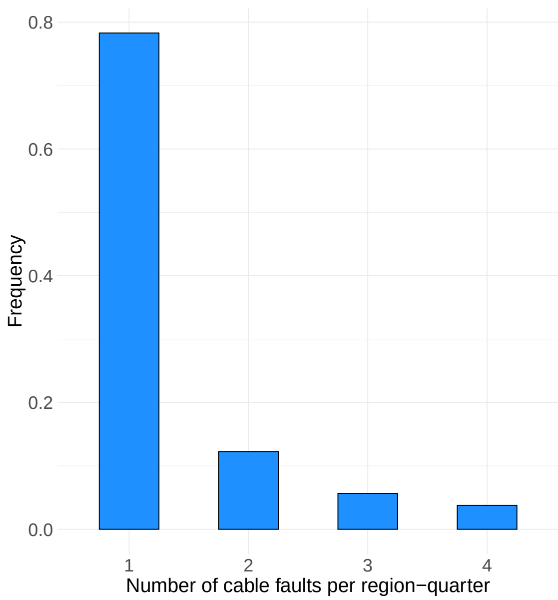

Figure A6: Cable Faults

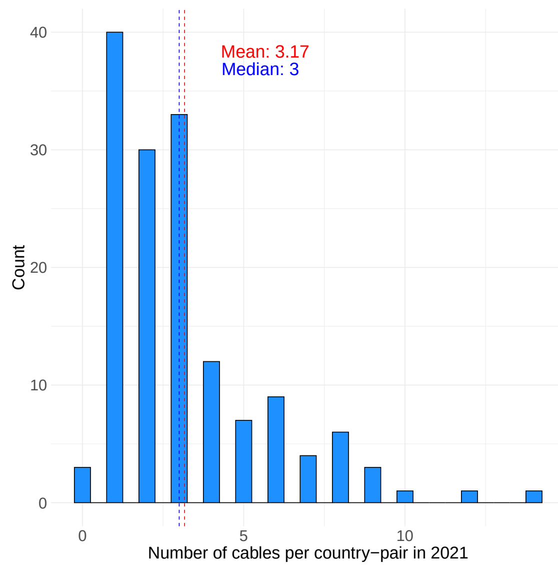

Figure A7: Cables

Notes: The left panel shows the relative frequency of different numbers of cable faults by region-quarter. The right panel shows the number of cables connecting country pairs as of 2021.

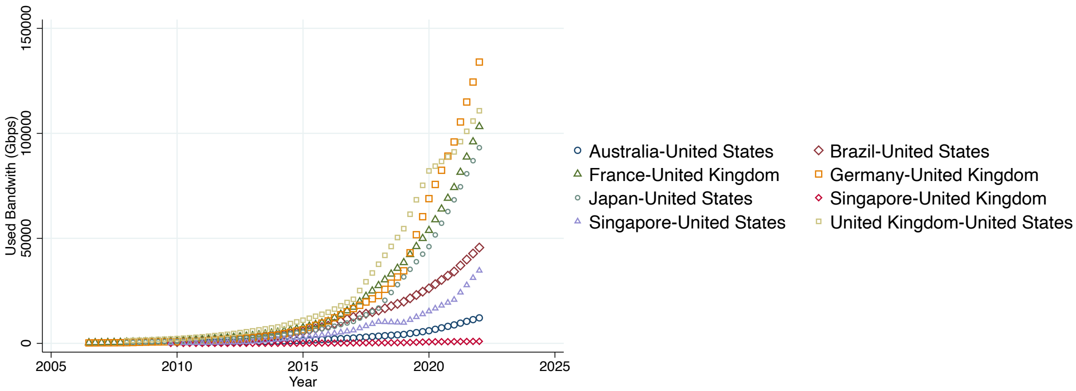

(a) Used bandwidth Q_{mt} in level

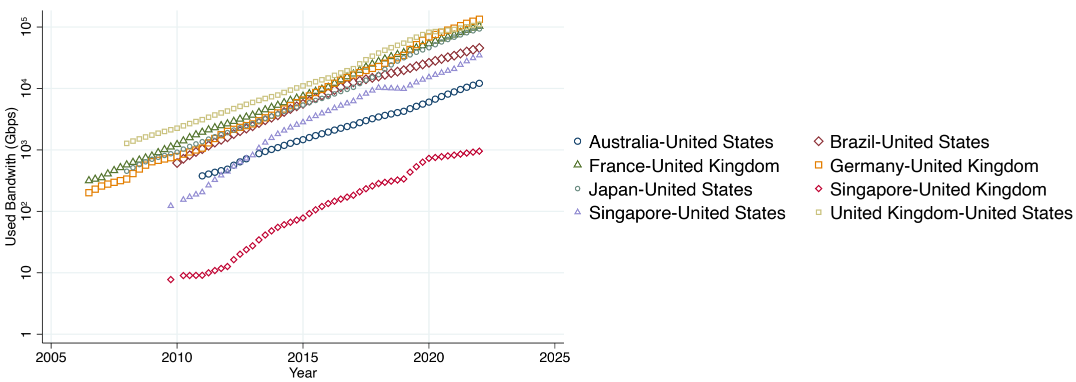

(b) Used bandwidth Q_{mt} in log scale

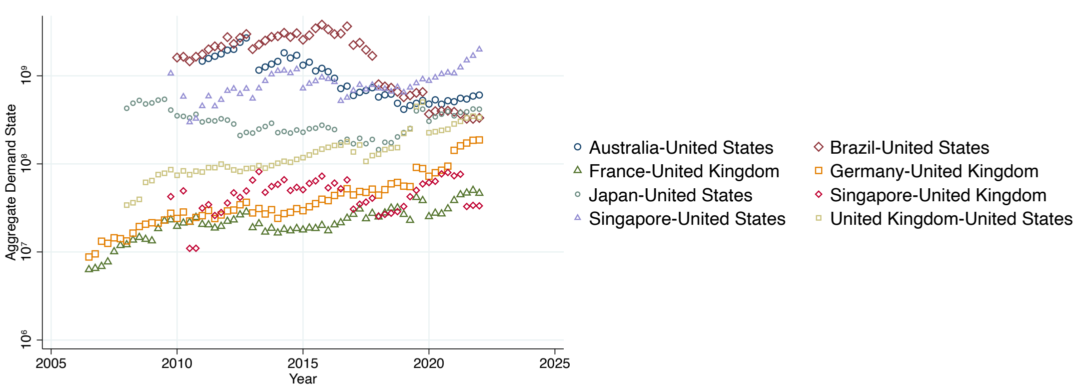

(c) Demand state d_{mt} (Gbps, log scale)

Figure A8: Used Bandwidth and Demand State Over Time for a Sample of Markets

Figure A9: Demand State Over Time for a Sample of Markets

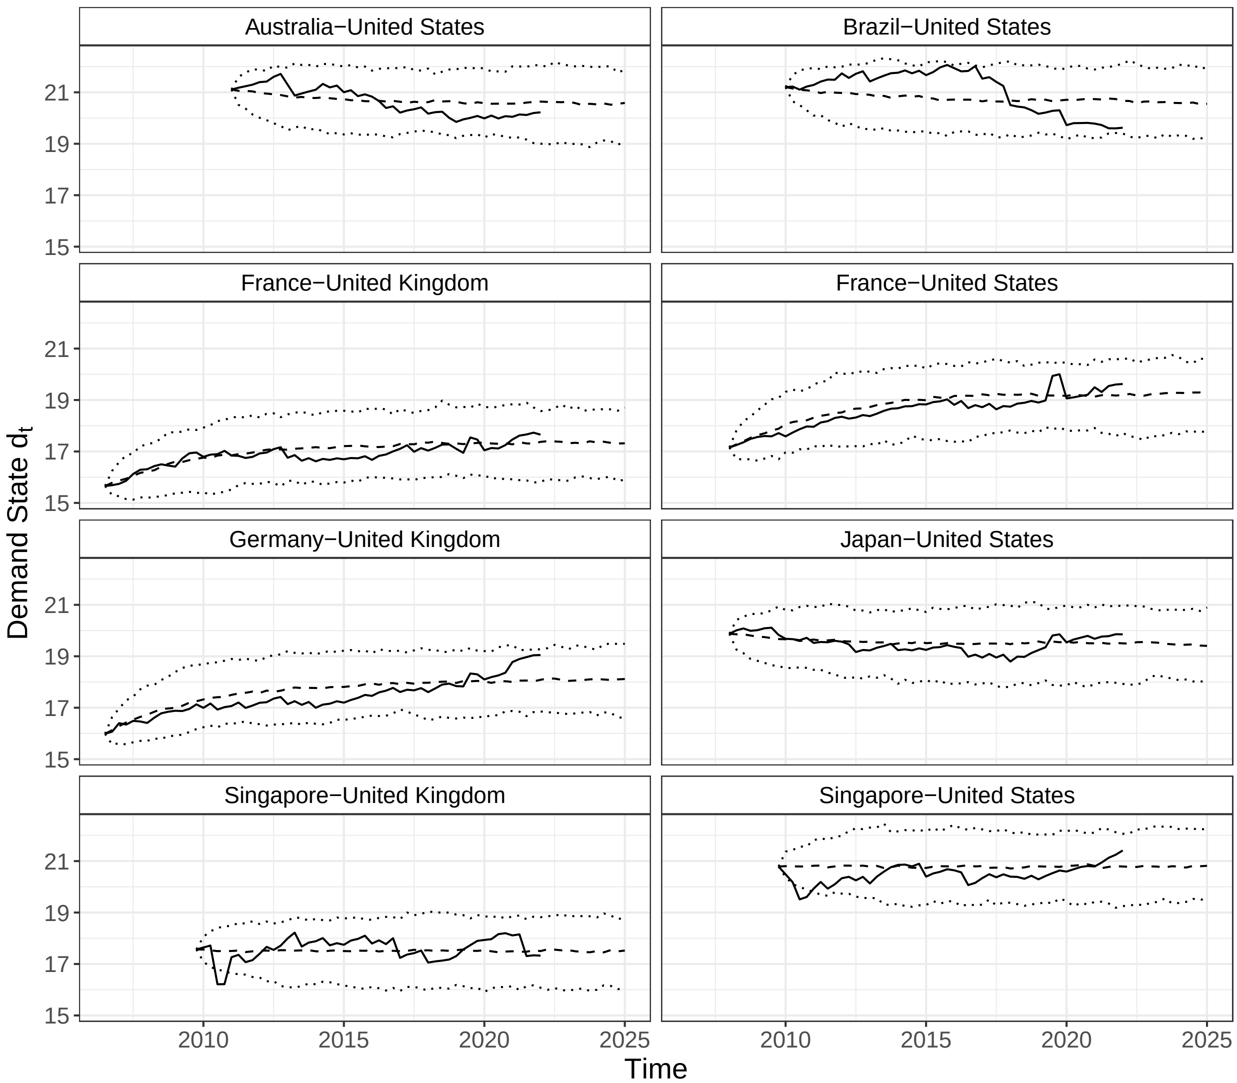

Notes: The figure shows the evolution of d_{mt} over time for a sample of markets (solid line). The dotted lines give the 95% confidence interval from simulations using the transition processes. The dashed line shows the median simulated value.

Figure A10: Cost State Over Time for a Sample of Markets

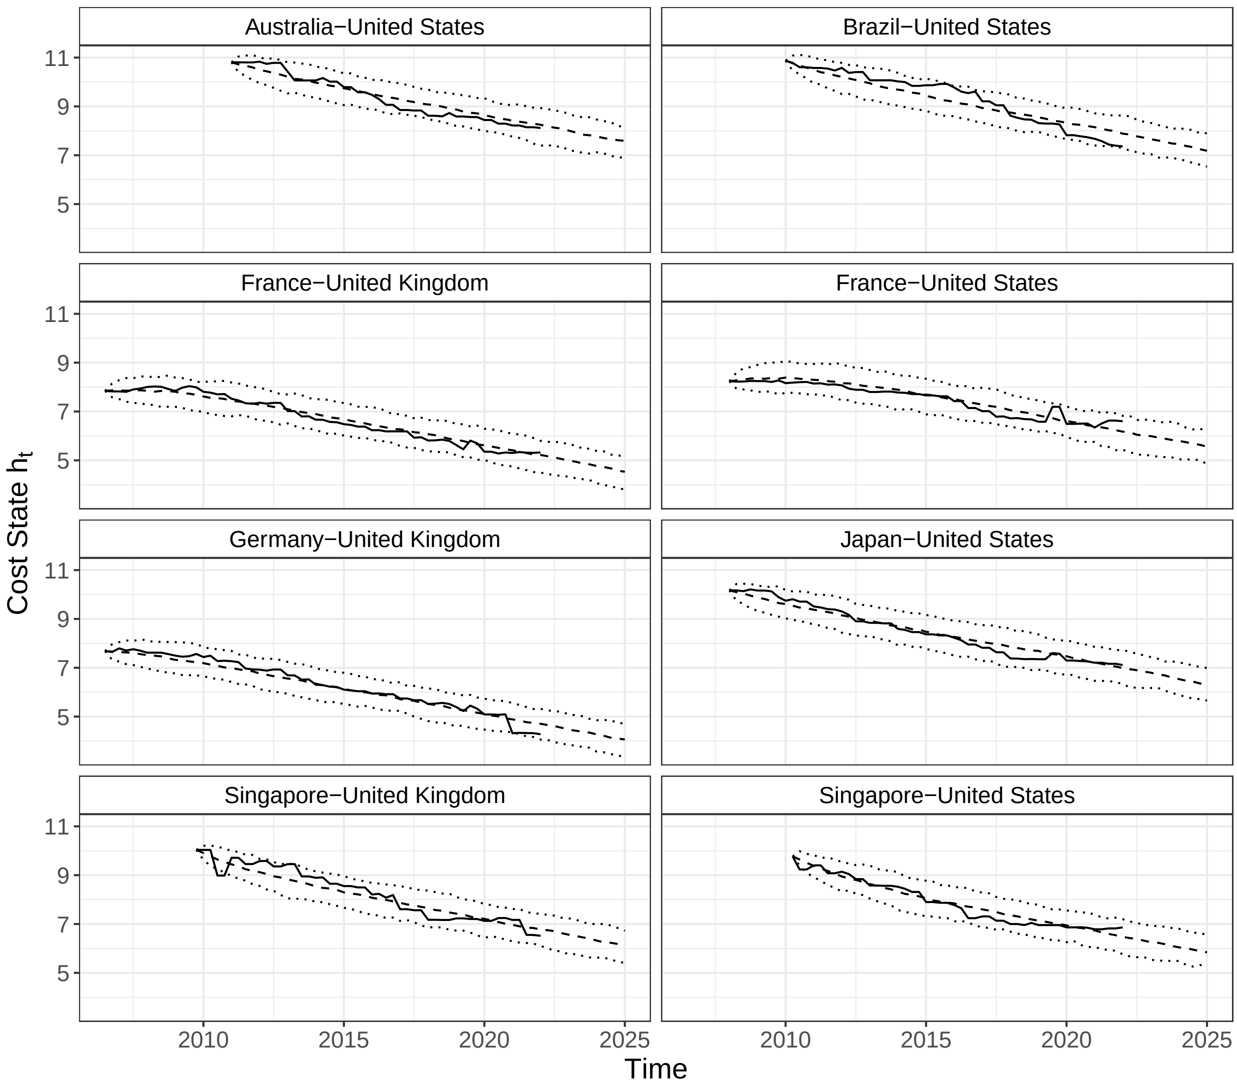

Notes: The figure shows the evolution of h_{mt} over time for a sample of markets (solid line). The dotted lines give the 95% confidence interval from simulations using the transition processes. The dashed line shows the median simulated value.

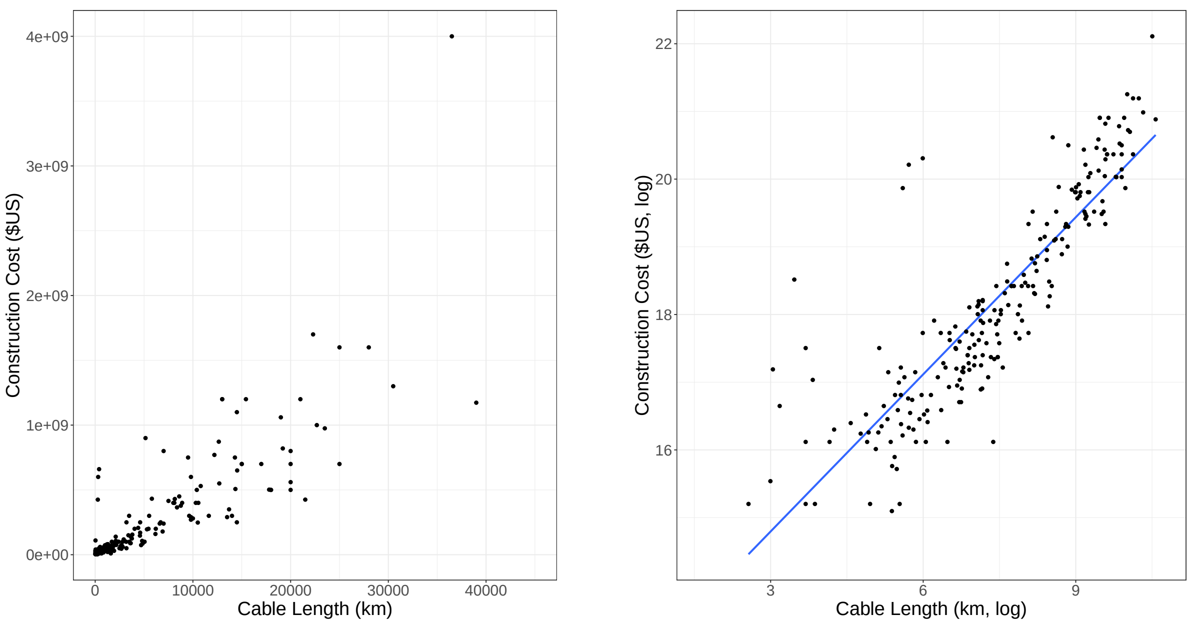

Figure A11: Cable Construction Costs and Cable Length in Level (left) and Logarithm (right).