

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Economic Analysis and Infrastructure Investment

Volume Authors/Editors: Edward L. Glaeser and James M. Poterba, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-80058-5 (cloth),
978-0-226-80061-5 (electronic)

Volume URL:

<https://www.nber.org/books-and-chapters/economic-analysis-and-infrastructure-investment>

Conference Date: November 15-16, 2019

Publication Date: November 2021

Chapter Title: Digital Infrastructure

Chapter Author(s): Shane Greenstein

Chapter URL:

<https://www.nber.org/books-and-chapters/economic-analysis-and-infrastructure-investment/digital-infrastructure>

Chapter pages in book: p. 409 – 447

Digital Infrastructure

Shane Greenstein

8.1 Introduction

The deployment and adoption of the commercial internet in the 1990s brought about a major restructuring of digital infrastructure. Today digital infrastructure supports a range of innovative businesses in the sharing economy, social media, mobile information services, electronic retailing, and ad-supported media. All such activities were much smaller in the 1990s, and their operations have changed dramatically in a few decades. These digital services continue to grow and take on more importance in GDP. For example, in 2017 electronic retailing reached more than \$545 billion for “Electronic Shopping and Mail Order Houses (NAICS 4541).” This category grew 65 percent from 2012. In 2017 online advertising contributed \$105.9 billion among “Internet Publishing and Broadcasting and Web Search Portals (NAICS 519130).” This category grew 250 percent from 2012.

While many economic studies focus on the most visible parts of the digital economy, much remains unexamined behind the surface at the level of digital infrastructure. This oversight neglects essential economic activity and overlooks the source of important productivity advances. The internet was

Shane Greenstein is the Martin Marshall Professor of Business Administration at Harvard Business School, cochair of the HBS Digital Initiative, and a research associate of the National Bureau of Economic Research.

Thanks go to numerous audiences, Klaus Ackermann, Andre Boik, Tim Bresnahan, David Clark, kc claffy, Chris Forman, Ed Glaeser, Avi Goldfarb, Bradley Huffaker, Maggie Kelleher, John Laprise, Kristina McElheran, Abhishek Nagaraj, Frank Nagle, Tommy Pan Fang, Jim Poterba, Jeff Prince, Bill Rogerson, Greg Rosston, Alicia Shems, Scott Stern, Catherine Tucker, Pai-Ling Yin, and Ran Zhuo for conversations, editing, and comments. All errors are mine alone. For acknowledgments, sources of research support, and disclosure of the author’s material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/economic-analysis-and-infrastructure-investment/digital-infrastructure>.




Fig. 8.1 The four layers of the internet

Source: Zhuo et al. (2019).

designed with a four-layer abstraction—application, transport, internet, and link (see figure 8.1)—so important parts of economic activity work with applications and are less visible. Processes in each layer handle requests from layers above and communicate with remote processes using layers below. This allows each layer to develop and operate independently. Thousands of independently owned, managed, and operated networks voluntarily exchange data via bilaterally negotiated agreements. These layers use many special types of equipment—root servers, fiber, broadband lines, networking switches and routers, content delivery networks, cloud facilities, and cellular towers. This equipment works with the privately owned investments of millions of content providers and user applications. Many organizations involved in digital infrastructure—such as data center operators, content delivery network specialists, data carriers, and access providers—perform this activity.

This chapter reviews studies from a number of areas, with an emphasis on innovation economics, industrial economics, growth accounting, and urban economics. The economic importance of these topics is almost self-evident. Even economists who are skeptical about public spending on traditional infrastructure still admit a role for large-scale public spending on digital infrastructure and on R&D to improve it. Among the salient questions examined in this chapter are the following: What determines variance in the supply of digital infrastructure, and how does that variance shape the performance of digital services? What does evidence suggest about the private incentives and economic returns to society from investment in broadband access and improvements in components of digital infrastructure? While statistical evidence documents considerable variance in the supply of digital infrastructure across regions of the United States, does that evidence show increasing or declining differences in the availability and use of the frontier over time, and why? Are the contributions of digital infrastructure to

economic growth correctly measured, and, if not, what are the lessons for measuring prices, output, and quality?

This chapter seeks to inform research and policy analysis, not to advocate choices over policy. That goal focuses the discussion and limits its scope. While the review informs questions about alternative proposals for regulating access, I do not take a position on a specific proposal or regulatory design of, for example, “net neutrality.” A curious reader can go to other sources for such analysis.¹ Relatedly, this review concentrates on economic research about the determinants of and returns from digital infrastructure in the United States and covers the global experience when possible. Again, a curious reader can go to other sources for international comparisons.² While the chapter assesses the economic justification for subsidizing infrastructure, once again, I refrain from performing an assessment of any specific proposal. Finally, the review does not provide a description of the minutiae of the engineering and operations of the internet. Again, other sources provide this.³

Section 8.2 discusses the legacy of the public origins of the internet and the adoption that followed after the internet privatized. Section 8.3 reviews the creation of value at homes and businesses and the role of innovation in creating value within the network. Section 8.4 reviews open questions about the governance of digital infrastructure. Section 8.5 concludes with observations about the unique boundaries between the public and private in today’s digital infrastructure.

8.2 Origins

The internet arose from combining multiple inventions, which were created with government support and operated by government organizations for public purposes. That experience left an imprint on the organization of the network. It exploded in its scale and scope after privatization in the mid-1990s as a result of investments on a massive scale.

At a high level, the architecture for the internet today bears some resemblance to its government-operated predecessor. One set of firms provides access, while another partially overlapping set of firms provides long-

1. This literature has focused on theoretical and legal issues, supported by examples, and not econometric measurement. See, for example, Nuechterlein and Weiser (2005) for a thorough review of the origins of many regulatory rules at the outset of the internet; Greenstein (2015) for a history of the growth of the commercial internet; and Greenstein, Peitz, and Valletti (2016) for a review of economic research on the topic.

2. A curious reader can start with OECD (2014), World Economic Forum (2016), and Cirera and Maloney (2019).

3. On the architecture and its evolution, see Clark (2018). On the basic operations of the transport and internet layer today, and the origins of congestion, see Clark et al. (2014). On the basic economics of data networks, see Greenstein (2020). On the operations of the network interconnection, see, for example, Norton (2014).

distance lines, and still other organizations provide a range of additional services, such as name-serving, domain registration, and routing. Exchange of data still follows parts of the models established in the late 1990s, albeit today at a much higher level and volume of traffic and with a different composition of firms and negotiating agreements. The level of interconnectivity is also much higher and continues to grow (Zhuo et al. 2019).

Numerous activities have evolved. The type of data today supports a different set of applications than in the 1990s. Email and file transfer once dominated; today, data-intensive applications are more prominent, such as video, streaming, and gaming (Huston 2017). Applications today accommodate a mobile user, and firms operate towers and antennae to support that use case.

8.2.1 The Origins of Internetworking

No simple model can describe the origins of the internet and why its architecture evolved as it did after its commercial applications grew. A brief overview of the stages of development prior to the transition into commercial markets can provide an outline of the complex changes:

- *Initial prototyping.* The first set of frontier inventions took place during the 1970s and 1980s when the Defense Advanced Research Projects Agency (DARPA) was the sole funder. Prototypes for packet switching were first engineered. The network at DARPA grew beyond prototypes, although the result was not technically straightforward at the time. The name, TCP/IP (Transmission Control Protocol/Internet Protocol), the specific design for protocols, became the label for this network. Contemporaries built much more around TCP/IP to make it viable.
- *Refinement of the network by the National Science Foundation (NSF).* In the mid-1980s, parts of the TCP/IP-based network were transferred to the NSF.⁴ Under NSF governance, the internet acquired a range of new refinements and new regional networks for supporting the shared use of resources.⁵ Further innovation took place in the domain name-server system (DNS). The Internet Engineering Task Force became established to guide further protocol development. These actions helped turn the internet into a living and evolving network, supported by a geographically dispersed organization.
- *Initiation of privatization.*⁶ During the early 1990s, the focus was prag-

4. The reason is that many civilian participants were frustrated by the challenges of getting military clearances and so on, and the NSF leadership foresaw benefits to the US academic research community. See Abbate (1999).

5. Until the NSFNET came into existence, there was only one network and one backbone, and BBN operated it. Eventually the NSFNET introduced additional backbones and regional carriers. See Abbate (1999).

6. These events are described in more detail in Greenstein (2015), chap. 3.

matic and oriented toward issues with scaling operations. The most important invention supported routing between multiple networks.⁷ A debate ensued about practices for data exchange in a privatized system in which, to achieve national interoperability of communications, competing firms had to cooperate. Independently, Tim Berners-Lee invented the World Wide Web and began to deploy it. It expanded the functionality of the internet in ways that made it more appealing to less technical users.

- *National deployment.*⁸ In the middle of the 1990s, the internet backbone became a private asset, allowing private firms to build on top of it. At the same time, at the National Center for Supercomputing Applications (NCSA) at the University of Illinois, a team developed the Mosaic Browser. It became the source for Netscape and Internet Explorer, catalyzing the “browser wars,” which occurred right after privatization. A related team at NCSA created a web server, which became the antecedent to Apache, the most popular web server for the next two decades.

To summarize, university researchers created many of these core inventions, and most received public funding from DARPA and NSF, with an exception for Tim Berners-Lee, whose funding came from CERN (the European Organization for Nuclear Research). The NSF and DARPA helped launch practices around standardization and network interconnection. Private markets inherited a reliable and operational network. While private firms had supplied some of the equipment used by NSF and DARPA, after the mid-1990s private investors picked up the bulk of operations and investment activity. While NSF-funded research into improvement in computing and networking continued after privatization, market forces took a more central role in determining the direction of innovation. The Web diffused on top of this infrastructure, and grew into the foundations for an enormous range of commercial applications.

8.2.2 Technological and Operational Legacy

At the outset, one key piece of the network infrastructure was visible to users, dial-up internet access. It built upon the existing telephone network, which was geographically ubiquitous prior to the diffusion of the internet as a result of public policies that encouraged universal availability of the phone network, even in high-cost areas. Dial-up services learned from existing Bulletin Board Services (BBSs), which provided experience with operat-

7. NSF switched from the routing protocol Exterior Gateway Protocol (EGP) and replaced it with Border Gate Protocol (BGP). BGP enables fully decentralized routing. Making this change was one of the early technical signs of the pending arrival of commercial network and the retirement of NSFNET. For extensive discussion, see Clark (2018).

8. These events are described in more detail in Greenstein (2015), chaps. 4, 5 and 6.

ing commercial services with dial-up technologies. The first generation of dial-up access was available almost everywhere in the United States within a few years after privatization.⁹

Not long after dial-up access demonstrated the viability of a national market, a variety of entrants aspired to provide faster speeds than delivered by dial-up. Today these access technologies go by the label “broadband.”¹⁰ Today, broadband providers dominate the supply of access services to the internet for both households and businesses. Broadband networks initially did not replicate the ubiquitous availability of dial-up. Broadband networks required physical investments. In low-density areas, the costs of building such lines were high, which necessitated high prices for access that some users were unwilling to pay. In addition, in recognition of a range of concerns, state and federal policy did not require universal geographic availability. Estimates of the population unserved by wireline broadband have varied from 15 percent to less than 5 percent of the US population, declining over time. Today about 10 percent of the US population does not use the internet, and surveys suggest the two prominent reasons for doing so are high prices and lack of availability (Anderson et al. 2019).

By historical standards, the switch from dial-up to broadband was swift. In 2001, only about one-half of US households had access to the internet, and virtually all access occurred over dial-up; today, approximately three-quarters of US households have broadband internet access in their homes. Most of that switch took less than a decade.¹¹ In its most common form, firms offer broadband as either DSL (digital subscriber line) over a phone line or through cable modems retrofitted to cable television systems. More recently, broadband over fiber has become available.

The experience today arose from three related diffusion trajectories. One, broadband access diffused to households. Two, broadband diffused to business, complemented by investment in advanced computing technologies. Three, different specialists offered activities, while a few large firms integrated into the infrastructure. Feedback loops moderated those trajectories. New application development encouraged more diffusion of broadband, which encouraged more application development. The uneven supply of

9. The reasons are discussed in detail in Greenstein (2015), particularly chaps. 5, 7, and 8.

10. The definition of broadband has undergone changes over time, as regulatory expectations change. For purposes of this discussion, the definition will be loose and encompass any wireline technology faster than the data rates of 56k dial-up, including ISDN, DSL, and cable modem service. Among wireless technologies, all Wi-Fi (IEEE 802.11b, g, n, and more); 3G, 4G, and 5G cellular services; and modern satellite service are broadband.

11. Greenstein and McDevitt (2011) provide data on the replacement of dial-up with broadband. For the diffusion of broadband, see Camille Ryan and Jamie M. Lewis, “Computer and Internet Use in the United States,” *American Community Survey Reports*, September 2017, <https://www.census.gov/content/dam/Census/library/publications/2017/acs/acs-37.pdf>; and Pew Research Center, “Internet/Broadband Fact Sheet,” June 12, 2019, <https://www.pewInternet.org/fact-sheet/Internet-broadband/>.

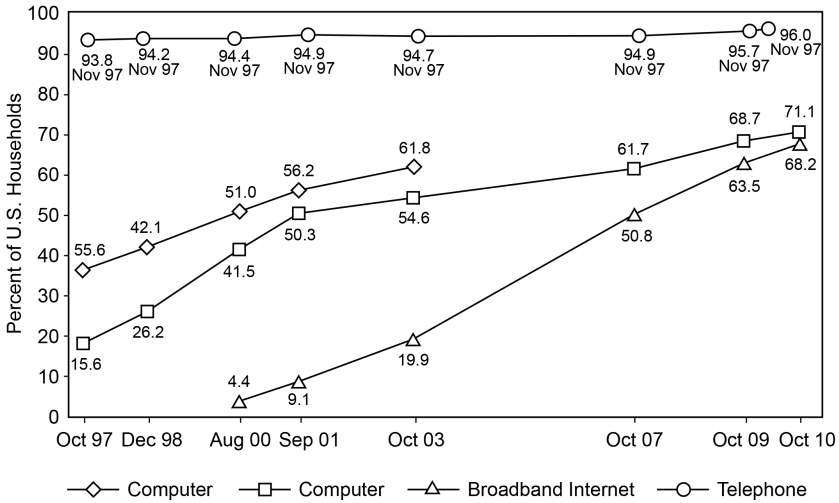


Fig. 8.2 Internet adoption over time

Source: National Telecommunications and Information Administration (2011).

networking created regional variance in the quantity and quality of digital infrastructure, and application development initially targeted areas and use cases favoring early adopters and, eventually, mass-market users. Uneven investment created variance in quantity and quality across countries, and different patterns of application development emerged within countries and across languages.

8.2.3 Adoption and Use

Adoption of the internet by households followed a standard S curve, with one-half of US households using dial-up by 2001 and one-half moving to broadband by 2007, continuing to grow thereafter. See figure 8.2, from a 2011 publication, for an illustration of both the S curves and the changing focus of US surveys.¹² At first they tracked personal computer use in homes, then dial-up, and finally broadband, ending the tracking of computer use.

Eventually, diffusion fell short of universal adoption as a result of lack of interest, lack of affordability, and lack of availability. Debates continue today over attributing nonadoption to different causes (see Horriagan 2020). Today, the adoption of broadband internet in the US hovers just below 80 percent of households, with the remainder of adopters using wireless

12. National Telecommunications and Information Administration, *Digital Nation: Expanding Internet Usage (NTIA Research Preview)*, February 17, 2011, <https://www.ntia.doc.gov/report/2011/digital-nation-expanding-Internet-usage-ntia-research-preview>.

access, either satellite or smart phones. In most developed countries, the percentage is higher. It is lower in underdeveloped countries, much lower in some. Worldwide, close to half of the global population uses the internet.¹³

The growing importance of digital infrastructure is also visible in other modes of access. For example, users of cellular telephony migrated from 3G to 4G, the latter entirely supporting digital communications.¹⁴ Presently, more than three-quarters of US households own at least one smartphone, rising from virtually none in 2007.¹⁵ Wi-Fi technology has diffused since its first deployment in 1999.¹⁶ More than 86 percent of homes with access to broadband employ Wi-Fi.¹⁷

What role does competition play? Wallsten and Mallahan (2013) ask whether competition plays an important role in improving quality to households. They examine the effect of more entry on the quality of the broadband provided, measured by its bandwidth, and then exclude suppliers from the count unless those suppliers provide services to a minimum threshold of customers. The authors found that—after considerable data cleaning—the typical zip code contained one or two suppliers of broadband, and a small number had three or more. This analysis shows that the third entrant does not change pricing but does generate competitive pressures for qualitative improvement.

This agenda extends in many directions. Seamans (2012) examines whether perceived threats of municipal entry generate faster upgrades and finds evidence that it does. Skiti (2019) examines whether potential competitive entry generates any response and finds evidence that it does. Chen and Savage (2011) focus on the role of competition in shaping pricing. They match cable and DSL internet access providers in all the western states and compare pricing differences between monopolies or duopolies in many small cities. The authors find that variety of customers mediates pricing and, generally, reduces price declines from an additional supplier.

Diffusion of the internet created two investment trajectories at business establishments. Forman (2005) proposed a framework for understanding adoption among a sample of early adopters, and Forman, Goldfarb, and Greenstein (2005) applied the framework to the US economy. They compare the use of *basic* with *advanced* internet technologies at US businesses near

13. World Bank statistics: <https://data.worldbank.org/indicator/IT.NET.USER.ZS?view=chart>.

14. 4G is the fourth generation of broadband cellular technology, succeeding 3G. 4G uses only packet-switching technology, unlike 3G, which used both packet-switching and (in parallel) the (old) circuit-switching technology. As of this writing, 5G contains much more capacity than 4G and has only just begun to deploy in developed countries.

15. Ryan and Lewis, “Computer and Internet Use in the United States.”

16. Wi-Fi is a standard defined by IEEE committee 802.11, operating over the 2.4 GHz and 5 GHz bands of spectrum.

17. NCTA, “Wi-Fi: How Broadband Households Experience the Internet,” April 6, 2018, <https://www.ncta.com/whats-new/wi-fi-how-broadband-households-experience-the-Internet>.

the end of the first wave of investment after the commercialization of the internet. Basic investment involved developing access to support email and browsing for employees, and a large fraction of establishments (approximately 90 percent) had adopted this. Advanced investment involved altering processes to supply services for customers and to receive inputs from suppliers, and a much lower percentage (approximately 12 percent) had adopted it. These latter activities were costly and depended on coordinating with partners and many complementary investments to enable electronic commerce (McElheran 2015).

Forman, Goldfarb, and Greenstein (2005) show that almost every establishment (approximately 90 percent) adopted the basic internet, while advanced internet showed up more prevalently in some cities. Several factors played a role in the deployment of advanced internet technologies. Some locations contained data-intensive industries that had recently made capital investments in computing and business equipment, which raised the returns to complementary investments in digital infrastructure. More educated and more skilled labor could take advantage of digital infrastructure, again raising the returns. Finally, some businesses were more productive and more profitable than other firms and, thus, could make bigger investments in all capital equipment, including digital infrastructure.

8.2.4 Innovation within the Network: Content Delivery Networks

Any improvement within the network improved performance for both households and businesses. The creation of content delivery networks (CDNs) provides an illustration of growth in specialists. CDNs first became available in the late 1990s and began spreading after that. Geographically distributed networks of servers located close to users, CDNs (i) reduce data delay by rerouting user requests and (ii) provide a layer of reliability and security.¹⁸ Today all but the smallest commercial content providers use CDNs. They have become an essential layer of digital infrastructure.

In the most common arrangement, a third-party commercial CDN negotiates with an internet service provider (ISP) or wireless access provider for the right to “colocate” a server close to users. The ISP may charge a “transit” fee to the CDN to take data over its network lines. Content providers pay the CDN to redistribute content to users from the CDN’s servers, which the content provider “updates” at an arranged schedule (by the minute, hour, or day). Many content providers choose to update only timely and popu-

18. Even when servers have gone down, the cached content in a CDN may keep a firm’s content available for users. In addition, CDNs can buffer content from a denial of service attack. For example, some CDNs, such as those operated by Cloudflare, have added security servers, such as protection against distributed denial of service (DDOS) attacks, which involve large numbers of queries to a server in a short time, exceeding the server’s capacity and rendering it unable to provide any service. CDNs are one of several instruments that can provide buffers against such attacks.

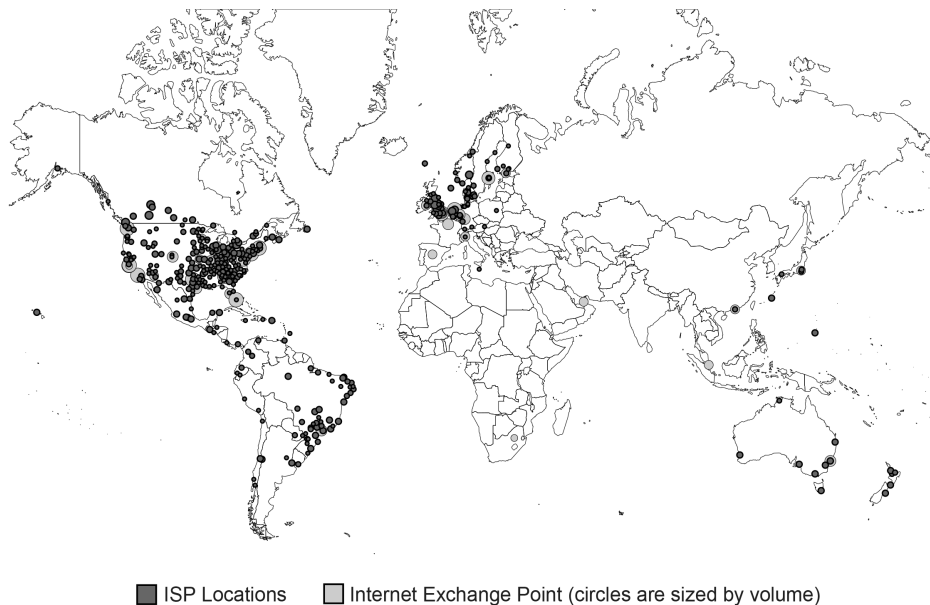


Fig. 8.3 Netflix CDN network

Source: Netflix Media Center, <https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience>, accessed February 2020.

lar content. Akamai is the largest provider of these CDN services in the US.¹⁹

Several large application firms have vertically integrated into operating their own CDNs. For example, Google, Apple, and Facebook operate CDNs and tailor the technical features of the CDN to their own needs. Again, they negotiate a price with ISPs for “colocation” in the network and sometimes pay fees for data transit. If negotiations with ISPs fail, the CDNs locate at internet exchange points (IXPs). In practice, only large firms opt for operating their own CDNs. It is usually less expensive to contract with a third-party CDN for small to medium volumes of traffic.

As illustration, figure 8.3 shows a map of Netflix’s CDN network. As one of the largest sources of streaming content in the world, the firm’s CDN network should be regarded as large. Netflix’s CDN network comprises more than 4600 servers in 233 locations in 2016, according to Bottger et al. (2018)—primarily deployed within ISPs or at IXPs in the developed world.

The growth of CDNs coincided with the improvements in consumer experience, in lowering latency for the large data flows supporting video.

19. Akamai’s revenues were \$2.7 billion in 2018. The next largest providers of such services, Cloudflare and Limelight, had \$192 million and \$184 million in revenue in 2018, respectively.

When the data packets traveled to users over dial-up in the mid-1990s, users typically could tolerate delays. Later, data traffic reached users primarily through broadband lines and became composed of mostly streaming, video, and gaming applications, but with fewer delays.²⁰ A symbiotic relationship emerged between improvement in access, CDNs, and applications. Many new applications would have been infeasible without CDNs, such as “over-the-top” streaming services like YouTube, Netflix, Disney+, or HBO Go.

The spread of CDNs frames questions about the economic impact of innovation. The gains distribute widely, while CDN providers make the investment. Content providers experience faster delivery, users enjoy previously unobtainable content, and ISPs charge colocation fees and gain revenue. Understanding these gains and externalities shed light on incentives to improve. The contractual arrangement involving third parties arises in virtually any but the smallest ISPs in the United States, which suggests the arrangement serves the interest of ISPs. In contrast, it is more difficult to infer that CDNs owned by content providers serve all parties’ interests. For a number of reasons—such as scaling issues, negotiating frictions, and the colocation expense—some firms prefer to locate some of their private CDNs at IXPs and not within ISPs. If application firms vertically integrate into their own CDNs and locate elsewhere, do others using third party CDNs get a different quality of service? As of this writing, no economic research has approached these questions.

8.2.5 Innovation: Data Centers and the Cloud

At the outset of the commercial internet, most firms housed their servers on company premises. That changed gradually (see, e.g., Byrne, Corrado, and Sichel 2018; Jin and McElheran 2018). Today third-party suppliers of data centers in the United States allocate assets worth at least several hundred billion dollars (Greenstein and Pan Fang 2019). Data centers lower latencies for business users, enable large-scale computing and innovative uses for that scale, consolidate managerial challenges and reap efficiencies from solutions to those challenges, enable flexible uses that previously were not possible, and remove frictions to accessing big-data applications. These abilities reduce the frictions supporting applications for a mobile labor force (see, e.g., DeStefano, Kneller, and Timmins 2019; Ewens, Nanda, and Rhodes-Kropf 2019). The largest agglomeration of data centers in North America is in Ashburn, Virginia, just outside Washington DC, near Metropolitan Area Exchange, East (commonly referred to as MAE-EAST), which is one of the oldest IXPs in the United States.

Just as with CDNs, the growth of data centers and the cloud illustrates an important question about the impact of investment in frontier digital

20. See, for example, the usage statistics in Nevo, Turner, and Williams (2016); McManus et al. (2018); and Huston (2017).

infrastructure. How do the gains distribute between cloud providers who operate the servers, the content providers who use them, and the users who enjoy previously unobtainable content?

Data centers contain rows of servers, which perform computation or storage. These buildings optimize for low energy use and optimal cooling, and they may contain expensive backup generators and structures to prevent flooding or reinforcements in floors to lower vibrations from passing vehicles. The inside wiring also may support a specific set of activities, especially in critical functions that support transactions with sensitive customer data.²¹ These expensive features can matter. For example, because of built-in resiliency and smart site selection, the data centers in Houston continued operating without interruption during and after the flooding of Hurricane Harvey in September 2017.

Contracts for data centers cover every conceivable arrangement and option between ownership and rental markets. At one extreme, rental markets arise for just about any arrangement a buyer could want. There are plenty of firms that will take responsibility for the operations of the building and electronic equipment for a service fee. Many buyers with generic needs—such as storage for backup—rent space in data centers at various time intervals (for example, 5, 10, or 20 years), own the servers and program them, and let others manage the building. At the other extreme, firms with unique computing needs, such as Facebook, Apple, Microsoft, Amazon, Oracle, and Google, own everything. They operate the largest private data centers in North America and configure the building and servers to suit their applications. For example, if a firm has an essential operation within a data center, large CDN networks typically complement it, so the results deploy quickly to users. Sophisticated firms increasingly utilize complex architecture to balance the loads from user demands, such as using CDNs for rapid response to requests for timely content, cloud facilities for secondary response, and remote servers for requests of the least popular content.

Today a cloud service involves a data center that rents its services, with the additional feature that users can request any size and turn the service off and on at will. The providers increasingly offer software services for a nominal charge or for none at all. Demand for cloud services has grown as the services improve in quality and decline in price. For example, Amazon Web Services offers scores of cloud software services; Microsoft Azure supports many Microsoft products, such as Outlook, as a cloud service; and Google offers TensorFlow, a standard tool for machine learning, at no charge with Google's cloud service. The appeal of the cloud comes from its flexibility in capital commitment and scale and the option to substitute variable

21. The data center for the New York Stock Exchange, for example, permits many firms to access trading services at especially fast rates. As another example, a segment of business users in health, finance, and transportation require high security and high reliability, often referred to as the “five nines” of reliability—namely, 99.999 percent uptime.

costs for fixed costs on a balance sheet, which appeals to cash-constrained firms.²²

The data center and cloud market have received some research attention. In the first paper on its productivity, Jin and McElheran (2018) examine the use of cloud computing in US manufacturing and find it predicts productivity growth among young firms and new units in established firms. Use of the cloud also predicts productivity, conditional on survival, in uncertain environments. The evidence is consistent with the highest gains accruing to firms that take advantage of the flexibility and lower costs of learning about IT needs in spite of uncertainty.

Tensions between the size of the investment and localization of demand shape the location of data centers. Greenstein and Pan Fang (2019) posit a framework that focuses on the “distaste for distance,” which creates localization of demand, and different supply conditions across geography.²³ Facilities spread out to match local demand. These compete with facilities that “aggregate” the demand from many locations. The costs of supply reflect variance in economies of scale and variance in operational costs. Both fixed and variable costs vary with the cost of inputs, such as land, electricity, cooling, and technical labor. These lead to variance in costs across different locations, and firms respond to this tension with entry and capacity decisions. Greenstein and Pan Fang (2019) forecast a “minimum threshold” of local users under which no entry occurs and find evidence consistent with this model. That suggests data centers and cloud services have an urban bias, favoring bigger and denser cities.

If buyers perceive shorter distances between users and the data centers for cloud services as an important attribute of cloud services, then that will create further potential for tension around the localization of supply. The first evidence about the demand for cloud services suggests users will place value on distance (Wang, LaRiviere, and Kannan 2019). While ubiquitous frontier infrastructure confers large societal benefits, such frontier infrastructure tends not to be available in low-density regions.

8.3 Creating Value

How and why did digital infrastructure produce value? How is that measured? The internet grew and diffused to households and businesses more rapidly than the telephone, electricity, and other technologies, so the ques-

22. See, for example, Coyle and Nguyen (2018). Byrne, Corrado, and Sichel (2018) estimate the quality-adjusted price decline between 2009 and 2016 at 17.3 percent per annum for Amazon Web Services.

23. “Distaste for distance” arises from a mix of three factors. The first two—user dislike for latency and user desire to avoid congestion—look alike in reducing distances between users and facilities. A third factor, “server hugging,” arises from managerial preferences for nearby physical facilities, which facilitates monitoring.

tion informs understanding of the causes of economic growth (see, e.g., Comin and Hobijn 2010).

8.3.1 Creating Value at Households

At first glance, internet access adoption would seem to follow the classic model of adoption, whereby those with the greatest willingness to pay adopt earliest and those with lower willingness to pay adopt later, as a result of declines in price, increases in quality, or both. In this model, the value to consumers provides the value of access in terms of consumer surplus. This model has considerable appeal because it provides a path toward valuing improvements from access infrastructure.

The model would appear to be a good approach for measurement. After all, just contrast the price and quality of internet access to households in 2001 with 2016. Around 2001, dial-up dominated access to the internet, and approximately half of US households were online. Web traffic dominated the internet, and wireless access had just entered a new era with the introduction of Wi-Fi and 3G cellular service, which ran a data service in parallel with voice services on cellular towers and handsets. By 2016, broadband access dominated all modes of access, and three-quarters of US households maintained connections online. In 2016, the predominant applications leading to data traffic were streaming, video, and gaming; Wi-Fi 5 and 4G served as the predominant wireless modes of transmission. This 15-year history suggests a large and valuable increase in access networks that should manifest in price declines, quantity increases, and qualitative improvement.

One positive symptom of improvement shows up in GDP (especially after the Census reclassified activities to track activity). From 2012 to 2017, payments for access to wireline forms of internet access reached \$88.7 billion, growing more than 30 percent in those five years. In addition, payments for access fees to wireless service reached over \$90.0 billion, an increase of 57 percent.²⁴

An estimate for user adoption by income, shown in figure 8.4, also seems to fit the model.²⁵ While adoption grows across all demographic groups, the variance in adoption across income is visible. The persistent pattern—with lower-income groups adopting less frequently—motivates hypotheses that high prices deter low-income households from purchasing internet access. Yet, other parts of the measured record contain more ambiguous indicators. The growth displayed in figure 8.4 ought to arise from either a decline in prices or an increase in quality or both. The consumer price index (CPI) for access covers only access. Proper accounting of user costs involves both a charge for telephone calls and a separate charge for internet access. For

24. Statistics of US Business, US Census.

25. These graphs aggregate periodic surveys (not smoothed) conducted by the Pew Internet and American Life Project.

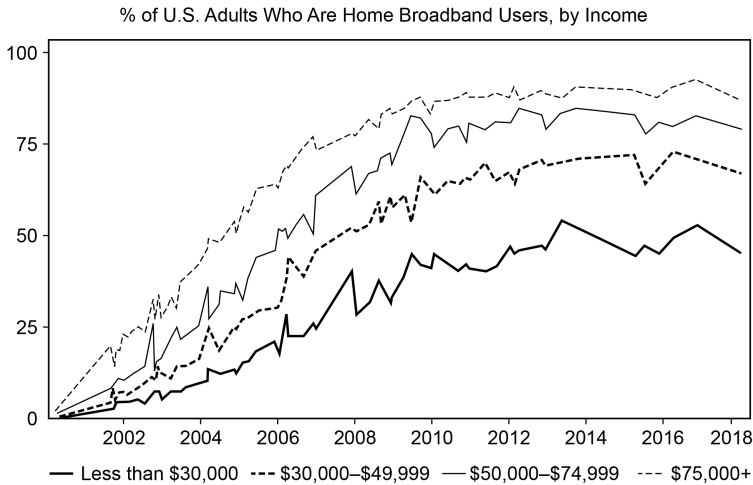


Fig. 8.4 Broadband adoption by income level

Source: See Pew Research Center, “Internet/Broadband Fact Sheet,” June 12, 2019, <https://www.pewresearch.org/internet/fact-sheet/Internet-broadband/>, under “Who Has Home Broadband?,” with income as the primary sorting variable.

some users the user cost included an additional expenditure for a second line. Many users sought to avoid the cost of an additional line, and those users employed the existing lines more intensely.²⁶

As broadband diffused, the CPI for internet access has covered broadband delivery, and the charges for telephony have become less relevant. That price series for internet access has remained flat for an extended period of time after a one-time drop in the middle of the decade. For example, in 2007, the CPI was at 73.2, and more than a decade later in 2018 it was at 76.0.²⁷ In other words, the consumer price of broadband has *increased* by 3.8 percent. The puzzle does not disappear with a comparison with other indices. The closest comparable CPI—that for wireless services, which also includes the price of telephone calls—displays a drop from 64 to 46 (a 28 percent decline in prices).²⁸

Simple alternative explanations do not provide an answer. Increased adoption cannot account for the rise in revenue in the face of no price change. From 2011 to 2018, approximately 3–5 percent of US households first began using broadband internet, depending on the survey. That is too

26. See Greenstein and McDevitt (2011) for the details.

27. See the series for internet services and electronic information providers in US city average, all urban consumers: US Bureau of Labor Statistics, “CPI for All Urban Consumers (CPI-U),” <https://data.bls.gov/PDQWeb/cu>.

28. See wireless telephone services in US city average, all urban consumers: US Bureau of Labor Statistics, “CPI for All Urban Consumers (CPI-U),” <https://data.bls.gov/PDQWeb/cu>.

small a number to account for a 30 percent growth in revenue.²⁹ Expenditure per household must have gone up, but how did that happen without a nominal price decline?

One explanation stresses that quality must have improved, but it went unmeasured. Some evidence suggests this is the case. For example, there is evidence of increasing speeds over time for all major wireline networks.³⁰ There are several potential reasons that the speed increase went unmeasured. First, as with many other consumer services, the CPI for broadband compares the prices of contracts for a given service.³¹ The procedure reduces measuring qualitative improvement if contracts do not reflect those improvements. That could happen because better caching, buffering, and other features do not factor into pricing in contracts. These features are hard to impute.

More subtle, contracts measure bandwidth and appear as part of a tiered menu of quality and price. If households do not use the same contract over the entire period, it is possible for them to increase expenditure on access without any measured price increase in a CPI. Hence, existing procedures also create an upward bias in the price index that fails to account for users switching to better contracts. That is exacerbated by the lack of measurement of savings, as noted earlier, when households dropped incurring expenditures for a phone line in order to move to broadband access.

A related explanation stresses issues with definitional boundaries between complements in use. The price measurement system treats access as a distinct service from content. Changes in the quality of content play no role in the price index for access. Stated another way, the standard measurement framework focuses on transactions for access, but not freely available services that users obtain along with their access. Standard procedures ignore the new goods—such as search, social media, and advertising-supported news and entertainment—even though content has improved. While those could motivate more adoption over time, as well as purchases of more bandwidth, we see only a constant price and more expenditure; not the cause, which is more quality.

There is secondary evidence to support these explanations. It draws from outside of price measurement and stresses the heavy evolution of applica-

29. More adoption did not produce the revenue increase. See Pew Research Center, “Internet/Broadband Fact Sheet,” June 12, 2019, <https://www.pewInternet.org/fact-sheet/Internet-broadband/>, and Ryan and Lewis, “Computer and Internet Use in the United States” (see n. 11).

30. See, for example, the Netflix comparisons of measured speeds over 2012–2018 yields a doubling of realized speeds for most networks (Netflix, “United States: Leaderboard,” <https://ispspeedindex.netflix.com/country/us/>, accessed April 2019).

31. The CPI is constructed from a weighted average of contracts for ostensibly similar services, where the weights come from household surveys and the contracts come from suppliers. This procedure necessarily underestimates the introduction of new goods—here, experienced as higher speeds—and qualitative improvements not reflected in common measures, such as bandwidth.

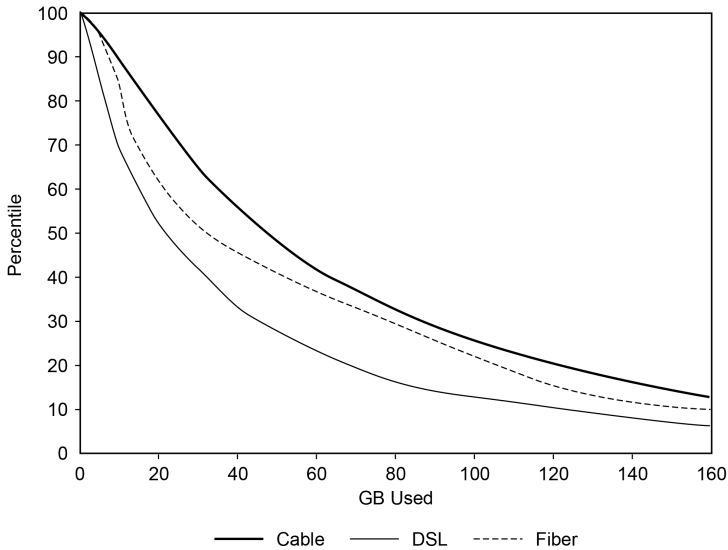


Fig. 8.5 Cumulative distribution of user traffic, by technology (highest users removed)

Source: Federal Communications Commission, September 2013, <https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-broadband-america-2014>, accessed March, 2020.

tions of the internet and the traffic that supports them. In the earliest days of the internet, text dominated traffic, either in the form of email or passive browsing. In contrast, more recently households have been adopting streaming services and receive increasingly many more magnitudes of data than they send (Huston 2017). For example, streaming a standard- to high-definition movie generates one to three gigabytes (GBs) per hour. To appreciate that size, examine figure 8.5, which shows a typical household's activity of data in 2013. The median household uses 20–60 GB a month. Netflix has increased its subscribership in the US from 20 to 60 million over the second decade of the millennium. Merely binge-watching a streamed series could massively increase data use. Netflix is far from the only streaming service. In short, as household streaming of television and movies rises, the capacity of access and underlying infrastructure also had to rise considerably. That could result in more intensive use of existing bandwidth and could motivate households to switch to more bandwidth at a higher price. That would register as more expenditure, not necessarily as a higher price in a price index.

Such considerations motivate research about heterogeneity in user demand for access. Rosston, Savage, and Waldman (2010) examine the demand for more speed as one of many attributes consumers choose to pay for. They find that a small set of users pay for higher speeds at any point in time. That

user (un)willingness to pay for more speed acts as a fundamental brake, slowing investment in upgrades in the short run. In other words, households act as if they prefer to migrate to higher speeds gradually, and firms respond accordingly.

Building on these estimates, Greenstein and McDevitt (2011) analyze the returns to households from upgrading to broadband from dial-up access. Despite the low valuations for frontier speeds, the authors show that the broadband upgrade over dial-up conferred a large consumer surplus on the economy. The consumer price index for access underestimated those gains, which would have generated at least a 2–3 percent decline in prices each year. Most conservative estimates of quality adjustment suggest an underestimate in standard economic measurement of access pricing.³²

Another informative research agenda analyzes both the contracts between users and access firms and subsequent user behavior. Usage-based pricing and data caps in wireline access contracts are more common, but only a little research has modeled the adoption and use decision in the presence of these contracting constraints. Nevo, Turner, and Williams (2016) provide a framework for understanding these decisions. They analyze usage data for a set of customers of a single ISP. These users face three-part tariffs, which impose a shadow value on the price of data as users approach their monthly allowances.³³ Users are sensitive to the charges affiliated with reaching a data cap, but they also endogenously select into capacity consistent with their own use, especially for those who are heavy users of streaming (as Nevo, Turner, and Williams observe, where 60 percent of data use relates to streaming applications). Variation in user behavior permits an analyst to recover variation in the willingness to pay for broadband, which provides insight into the gaps between private and social incentives to build or upgrade broadband. The estimates of Nevo, Turner, and Williams (2016) suggest that the gap is substantial, once again consistent with the presence of insufficient private incentives to upgrade quality at a rate in line with society's broader interest.³⁴

Byrne and Corrado (2019) focus on valuing the missing free complements. Borrowing insights from the measurement of capacity utilization, the authors argue that some consumers use access technologies for free goods more intensively than others.³⁵ The complementarity between paid

32. It is important to note an additional implication. The gains should result in a reallocation of household time, which will generate restructuring in many other industries that also bid for household time, such as television, radio, news, and entertainment.

33. See Burnham et al. (2013) for an early census of the use of tiered pricing and caps based on the usage of data in wireline and wireless forms.

34. Malone, Nevo, and Williams (2017) also examine the willingness to pay for more bandwidth, based on usage data from one ISP. They focus on the trade-offs for different ways to approach congestion of networks. The authors show that peak load pricing along with caching more effectively deals with congestion than does throttling of traffic.

35. This approach follows numerous studies that examine the time spent online as a possible avenue for valuing digital goods. See, for example, Boik, Greenstein, and Prince (2019); Bryn-

access services and network use leads to an unmeasured quality adjustment for the price of access. Looking across cable television, cellular telephony, and the internet, the authors calculate a nearly \$1,800 boost to consumer surplus per connected user, which amounts to a one-half percentage point addition to US real GDP for 2007–2017. That suggests the derived demand for access infrastructure is large, and so is its underlying value. None of the free services could provide such satisfaction without employing digital access and relying on nearly ubiquitous access.

The progress in understanding the experience outside the US has tended to take advantage of idiosyncratic institutional details that create opportunities for insights. An early paper comparing international experiences is Wallsten and Riso (2010), which shows a wide variance in access prices and availability across countries. Yet that has not stopped naive approaches that reduce the nuances in prices to one statistic to facilitate comparisons.³⁶ Wallsten and Riso's findings suggest that single statistics heighten the potential for unobservable factors in cross-country regressions.

One set of studies examines the deployment of broadband in the United Kingdom. As a result of the UK's underlying switch network, broadband deployed in a somewhat random geographic pattern, creating similar neighboring areas with different broadband experiences. This quasi-randomness creates plausible exogeneity. One line of research looks at the consequences of uneven broadband deployment on property prices for homes (Ahlfeldt, Koutroumpis, and Valletti 2017). Better broadband has an impact on local prices for real estate, evidence that homebuyers value broadband.

8.3.2 Creating Value at Business

The experience within business in the 1990s creates challenges for measurement. The deployment of email and browsing cannot generate insight into whether adoption of novel digital technologies had an impact because these basic internet technologies became available *and adopted* almost everywhere in the US within a few years, leaving almost no variance from which to infer the gains. At best, this diffusion would show up in general gains in productivity, though growth accounting would not be able to attribute the growth to any specific investment.

What can be inferred? It is possible to examine changes consistent with the adoption of advanced internet technologies, which required broadband and complementary investments as well as skilled labor, and for which there is variance in supply across regions and industries. This is the approach in Forman, Goldfarb, and Greenstein (2012), which considers whether the investment in advanced internet technologies became associated with alleviating

jolfsson, Collis, and Eggers (2019); Brynjolfsson and Oh (2012); Goldfarb and Prince (2008); Goolsbee and Klenow (2006); and Hitt and Tambe (2007).

36. Most commonly used are OECD (2014) or World Economic Forum (2016), which get their broadband prices from the same source: data from the World Bank.

or exacerbating regional inequality. Building on research linking information technology use to productivity gains,³⁷ an optimistic view forecasts that digital infrastructure potentially could reduce distances and aid those who lived at a distance from areas with higher incomes. The authors find, in contrast, that the first wave of the investment boom exacerbated regional inequality. Using an instrumental variable approach and a battery of additional tests, the authors relate wage growth to investment in advanced internet technologies. They find that business adoption of the internet makes regions with higher income richer in some places, but not everywhere. The largest divergences occur in major urban areas with skilled workforces and prior investment in IT. In short, the high-income locations experienced the most wage growth.

Note the large open question: Has subsequent regional growth altered the pattern of additional investment in digital infrastructure by business? Has additional investment continued to produce wage dispersion? Have regional IT wages diverged from other skilled wages, or have skilled wages diverged in a similar pattern? Relatedly, why has IT investment continued apace while real productivity gains have stayed close to 2 percent per annum after the much higher rates of productivity growth during the dot-com boom? How has that productivity growth been distributed across the country, and does it bear any relationship to the regional variance in the first generation of advanced internet equipment?

Inadequacies in data also make it challenging to infer the productivity effects of broadband. A researcher typically has access only to either (i) available supply of broadband or (ii) purchased supply of broadband. Each suffers from a distinct form of endogeneity bias and measurement error. There are additional challenges to measurement. At one time, the Federal Communication Commission (FCC) ostensibly tracked the former at the geographic level of the zip code but counted any firm as a supplier if it had one customer in that zip code. By including satellite suppliers, the FCC came to numbers that reached maximal levels. At best, less availability—when measured as zero or one supplier—indicates a setting with limited supply and little competitive pressure. It is challenging to find an econometric escape from such limited data. Today the federal government provides a broadband map of availability; more availability does not tell us about adoption or use.³⁸

One approach to these challenges, taken by Kolko (2012), examines different indicators of economic change affiliated with broadband—growth in information industries, wages, employment, telecommuting, and home-based work—and focuses the investigation on medium-sized cities where

37. For a recent review, see, for example, Cardona, Kretschmer, and Strobel (2013).

38. Federal Communications Commission, “Fixed Broadband Deployment,” <https://broadbandmap.fcc.gov/#/>.

exogenous instruments might be plausible, such as the topography of an area. This approach does not lean heavily on any single finding, because of suspected measurement error. The approach looks for robust patterns. Kolko finds that many indicators of a relationship between more broadband and improved economic activity, though not the key ones affiliated with taxes, such as wages and employment.³⁹

The experience outside the US has some parallels but also generates new insight. DeStefano, Kneller, and Timmins (2018) take advantage of the uneven rollout of DSL in the UK and link that to information about firm productivity. They find that the impact of broadband on business productivity is modest at best. They do see, however, that broadband is associated with restructuring the location and scale of activity. These results suggest complementary investments can play a significant role in fostering restructuring organizations, even when no short-term productivity improvement is visible.⁴⁰

Using detailed information about wages and workers, Poliquin (2020) finds another parallel experience with business adoption of broadband at Brazilian firms, where broadband became geographically available in a quasi-random way to firms. Overall, wages increased 2.3 percent on average at establishments following the establishment's adoption of broadband, consistent with a productivity gain. Consistent with the theory of biased technological change, wages increased the most for workers engaged in nonroutine cognitive tasks, while returns were negative for routine cognitive tasks. There was no effect of broadband adoption on wages for either routine or non-routine manual tasks. Poliquin also finds skill bias arises from changes within an existing labor force, not additions to it through recruiting.

The economic impact of access extends to topics around the globe. For example, one set of studies examines the impact the global spread of digital infrastructure spread had on trade, such as Fernandes et al. (2019). They examine export behavior in China during the period 1999–2007 and link firm participation in export markets to the rollout of the internet. They combine firm-level production data with province-level data on internet availability. Manufacturing rose during this period, and they find evidence of improvements in communication with buyers and input suppliers coincident with a more visible virtual presence. Like other studies, this one finds that improvements depend on the availability of broadband, but broadband alone is insufficient to explain all the increases in manufacturing. The authors stress

39. In spite of concerns about measurement, this continues to be a popular approach for measuring availability, particularly on the margin for lack of availability. For example, see Falck, Gold, and Heblich (2014).

40. This is in line with other work on the impact of broadband on the productivity of business, which also find modest effects on productivity but measurable changes in other firm attributes in the presence of complementary investments (see also DeStefano, Kneller, and Timmins 2019; Haller and Lyons 2015, 2019).

the role of numerous complementary investments to implement productive uses for broadband in firm processes.

Deployment of broadband access also generates symptoms of economic growth, especially in locations that previously lacked any wireline access. Hjort and Poulsen (2019) examine the gradual laying of fiber along the African coast, which enables wireline access where it previously was only possible by satellite. This experience is exogenous to potential adopters and creates many points of comparison between the served and underserved areas, including improved adoption (see also Cariolle 2019). Hjort and Poulsen's estimates show large positive effects on employment rates, with especially large increases in high-skill occupations. Remarkably, the authors also find smaller gains in employment for less educated workers.

An important open topic concerns the effect of digital infrastructure on entrepreneurship. In the developed economies, digital infrastructure has played a role in fostering a test bed for frontier application development by a range of entrepreneurial business. The size and importance of such externalities remain elusive to quantitative methods. Comparison of worldwide experience holds potential to identify the role of infrastructure in fostering technology-led entrepreneurial effects.

8.3.3 Value from Improving Protocols

A key piece of network infrastructure is its protocols and protocol stacks, intended to make digital equipment universally compatible.⁴¹ Although complex, the protocol stack design for sending data packets along the least congested route was, and continues to be, an essential feature of the digital infrastructure inherited from the NSF/DARPA era. Many improvements continue to be added to this design.

Protocol development today does not reside exclusively with governments. Several nonprofit organizations design and upgrade the protocol stack used for global internet infrastructure. Many stakeholders contribute to improvements. For example, the Internet Society oversees the Internet Engineering Task Force (IETF), which designs protocols behind TCP/IP (Transmission Control Protocol/Internet Protocol), BGP (Border Gateway Protocol), and other protocol stacks. The Internet Society and other organizations subsequently charge little for their use.⁴²

An improvement in the protocols of the internet benefits households,

41. Protocols are the set of rules and regulations that determine how data makes it through the network. A networking protocol defines conventions for processes, which includes definitions for both the format of data packets and also for recovery in the event of transmission errors. Protocol stacks are composed of a family of related protocols assembled together; they act as a reference model for designers, who largely aspire to make compatible equipment. For longer descriptions, see, for example, Clark (2018), Greenstein (2015), or Knieps and Bauer (2016).

42. The Institute of Electronic and Electrical Engineers (IEEE) maintains 802.11, the standard underlying Wi-Fi. The Internet Corporation for Assigned Names and Numbers (ICANN)

businesses, carriers, and application developers. Nonrivalry in use, combined with de facto lack of excludability, gives software protocols a set of properties isomorphic to classic public goods. This implies that improvements in protocols could confer large gains to society, and failures to design well could have negative consequences. Hence, similar to the questions for CDNs and cloud computing, developments in protocols raise questions about estimating externalities in an interdependent system.

The estimation of these gains or consequences is challenging because every user and supplier has access to the same improved protocols at the same time. Relatedly, as is true for many public goods, there is no competitive alternative, and users cannot opt out of the protocol stack if the relevant institutions make poor decisions. That (typically) results in no meaningful variance in adoption with which to make estimates of impact.

An important example of research about protocols is Simcoe (2012). The author examined the speed with which the IETF generated new protocols for the internet before and after privatization. He traces variance in speed to its underlying determinants, such as the composition of the committees making new protocols. He focuses on the role of disagreements between participants with varying interests, stressing the importance of multistakeholder institutions that (do or do not) become slower as the private costs and commercial risks become concentrated in the efforts of a few experts, who either come from universities or firms. Their interests may not align as their designs touch a wider breadth of the economy, resulting in a greater breadth of voices developing conflicting stakes in the details.⁴³

Other research focuses on the growth of standardized and large-scale (and mostly invisible) digital processes for collecting and reselling a user's data and for supporting auctions to place online advertisements (Goldberg, Johnson, and Shriver 2019). Market incentives have not produced a system that transparently informs users about which aspects of their private data will be collected and sold to others after use of an application. Obtuse terms of service have proliferated, and every user must possess nearly an advanced degree in computer science and legal scholarship to figure out how to answer basic questions about whether their data will be resold. Attempts to design a standard infrastructure for privacy, such as P3P, failed to be adopted (Cranor et al. 2008). It is a remarkable state for a feature with such public importance.

That example illustrates an open avenue for research. What are the incentives for and gains from improvement in protocols, as designed by quasi-

governs assignment of domain names and updates the routing tables used by every switch and router on the internet. See, for example, Clark (2018).

43. Research with this focus has also begun to explore the workings of committees who govern other important elements of equipment. For a variety of perspectives about standards in wireless communications, see, for example, Bar and Leiponen (2014) or Baron, Gupta, and Roberts (2018).

public organizations or private firms? How do these incentives align with the incentives to adopt such protocols? How large are the shared benefits of improving protocols? More broadly, today a mix of publicly subsidized and privately funded research finances protocol development. How large are the contributions from members in relation to those benefits? Government users no longer act as the major test bed for protocol development as they did in the past. Whose experience has the most salience for the direction of improvement?

One approach to these questions focuses on estimating the size of the externalities from the deployment of a key piece of infrastructure. For example, Nagle (2019) takes a novel approach to this topic by examining a set of externalities in protocol improvement that *were not* global. Counterintuitively, he focuses on quite the opposite: externalities from software in which the spillovers were particularly *localized* in scope. He focuses on the spillovers from a French government mandate to use Linux, a program adopted as part of general policies to encourage the use of open-source software. Nagle finds the mandate had consequences for the rate of new business formation in complementary digital areas. The analysis takes advantage of a natural placebo test in events, in which the Italian government did not enforce a similar decree within its own borders. Nagle's estimates suggest that the externalities can be substantial if governments enforce their policies. The study frames a big open question: What conditions lead the local supply of talent to respond, and what limits that response?

Another approach to this topic examines an episode of the recent past, in which, with the benefit of hindsight, the economics were comparatively simple. The costs of R&D were defrayed against the benefits affiliated with meeting the mission of a federal agency (at DARPA and NSF), and professional recognition among research peers provided motivation for most of the efforts. While these costs were concentrated, the external benefits to society were widely shared. That sets up a question: What were the economic gains from the public investments in the historical R&D that supported protocol development? Greenstein and Nagle (2014) employ a method for estimating the value of unmeasured web servers in the United States in 2011. The authors show that these inputs make a positive contribution to economic growth in the United States. The authors further show that the returns from web servers alone generated enough economic gains to equal the US government's R&D internet investment. That is an important conclusion, because the authors do not make a full account for all gains from the invention of the internet (which is still an open question).

A third approach analyzes evolving market-based events using economic lessons from outside digital infrastructure. For example, the exhaustion of the Internet Protocol address space markets for trading IP addresses. Edelman and Schwartz (2015) consider alternative principles for organizing the

design for the new and resale market and the properties affiliated with different proposals.

8.3.4 Uneven Geographic Deployment

Much research focuses on understanding the causes behind and consequences from the uneven supply of access and related network components. In developed countries, users in most suburban locations saw many options, with occasional overbuilding leading to more. Businesses in high-density settings could experience even more supply (see, e.g., Connolly and Prieger 2013; Wallsten and Mallahan 2013). Beyond those simple statements, the actual experience with entry and adoption depended on a host of factors, such as regulatory rules for pole attachments and ease of interconnection.

Variance in supply potentially creates the type of variance that econometricians like to exploit. The primary challenges are measurement. Many of these variations cannot be seen except at a fine level of geography, such as a neighborhood. Attempts to measure availability at this fine-grained level have encountered numerous challenges. For example, an attempt to create a National Broadband Map began in 2011, went through several revisions, was regarded as accurate in some but not all locations, and was discontinued in December 2018. As of this writing, the FCC is developing a new mapping program.

Government subsidies for high-speed access networks arise partly from analogies with local telephony, in which many providers received building and operational subsidies from universal service programs.⁴⁴ For example, the 1996 Telecommunications Act established the E-Rate program, which taxed telephone calls to finance subsidies for rural broadband. Today the program raises more than \$4 billion dollars annually, focusing on developing broadband internet access in costly locations and making it available to organizations with public missions, such as libraries, schools, and hospitals. As another example, the 2009 stimulus package included \$7 billion of subsidies for rural broadband. At a local level, many local governments also try to shape supply. Many insist through cable franchise agreements that cable providers build out into low-income or less dense areas.

Programs to address demand also exist but are less common. With this motivation, Rosston and Wallsten (2020) examine the impact of the Internet Essentials program, sponsored by Comcast to foster adoption of broadband by lowering prices for qualified low-income households in the parts of the US where Comcast provides service. This program provides information to test the proposition that low-income households are reluctant to adopt

44. There were many proposals for rural subsidies of broadband as part of the 2009 stimulus package, and they built on a previous set of subsidies in the E-Rate program, which were established by the 1996 Telecommunications Act.

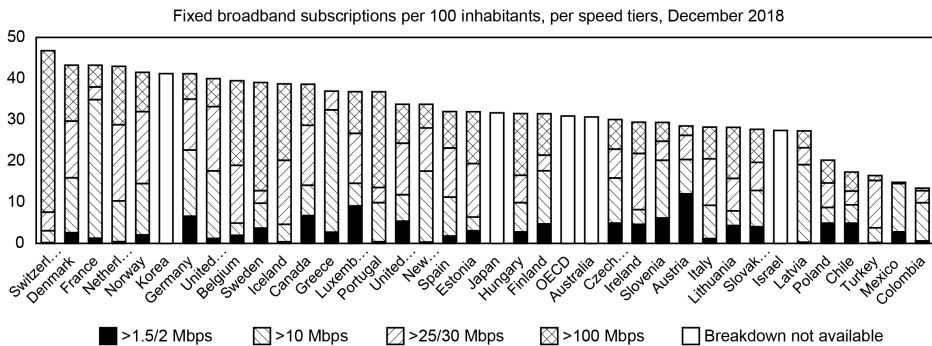


Fig. 8.6 Broadband across the globe

because of high prices. The measurement challenge requires clever statistical approaches. Rosston and Wallsten estimate *new* adoption in comparison to the counterfactual—that is, some households who qualify for the program would have adopted at the regular higher price. By comparison with adoption among similar households in similar areas that lack such programs, demand does grow among the target population in areas where the Internet Essentials program operates, suggesting the program supports hundreds of thousands of users who would not have adopted otherwise.⁴⁵

Another type of study focuses on rural broadband (see, e.g., Whitacre, Gallardo, and Strover 2014). An interesting fact complicates inference: broadband satellite has been available in virtually every location, and for many years. For many uses, such as email, browsing, and noninteractive internet services, satellite broadband is technically sufficient, albeit more expensive than broadband in a typical suburban location. With such facts as motivation, Boik (2017) investigates a specific situation to understand the micro-mechanisms shaping behavior. He examines low-density North Carolina and studies willingness to pay for satellite broadband versus wireline broadband. He finds considerable willingness to pay for wireless access, which, in turn, limits the potential welfare gains from subsidies for building out wireline access. This willingness renders uneconomic most subsidies for wireline services in low-density locations.

Many open questions remain. Considerable data exists to measure variance across the globe (OECD 2014; World Economic Forum 2016). Figure 8.6 illustrates broadband capabilities across the globe. Why does this variance arise? What economic outcome does the variance produce in dif-

45. Rosston and Wallsten also show that a fraction of current nonadopters (potential users) are insensitive to price—in the sense that a large number of nonadopters do not change their behavior in spite of these massive price reductions. This suggests nonadoption among laggards does not have economic causes and requires policies not focused on price.

ferent countries? Studies of micro-mechanisms could illuminate the causes and consequences.

There is also need for analysis of the experience outside of developed economies. For example, Björkegren (2019) examines the demand for and benefits from mobile digital infrastructure in Rwanda. Using a year of phone calls, he provides estimates of the value of belonging to a network, as well as of the value of the infrastructure that supports it. Here he finds evidence of network effects in demand, which suggests the externalities from new networks can be substantial.

Further estimates of demand for wireless access in developing countries are needed. Unlike other innovative products in the modern economy, all innovative digital services do not first arise in developed countries before migrating to the markets of developing countries. A set of innovative services—and new to the world!—have begun to appear in the developed world where wireless devices have become the primary tool for accessing the internet. Many fundamental economic activities, such as payments and banking, have developed atop this ubiquitous wireless infrastructure. For example, China’s most popular payments application for wireless devices, WeChat, has more than one billion users and has become an electronic substitute for cash. Such innovation has become increasingly common in developing economies and merits further study.

8.4 Open Research Questions

Two distinct views animate open questions about digital infrastructure supply. An outlook that could be labeled as “optimistic” anticipates experimentation in a few places, followed by more diffusion to more users, more regions, and a larger set of applications. This view interprets the state of digital infrastructure at a point in time as temporary, transient, and in the midst of wider diffusion. In contrast, an outlook that might be labeled as “pessimistic” stresses that digital infrastructure has achieved higher productivity in dense locations. That arises because of economies of scale in equipment, increased productivity from the colocation of many related activities, and the availability of skilled labor in urban areas in developed economies.

The two outlooks make different predictions and, accordingly, base policy on different premises. In the optimistic outlook, differences in supply melt over time as once expensive infrastructure, which incubated in a few cities, spreads to new users and new locations. The most important open question concerns the determinants of the speed of diffusion, which then determines how fast laggard regions catch up to frontier regions. In this view, policy intervention focuses on speeding up diffusion to laggard locations by removing deterrence to adoption. The pessimistic outlook perceives persistent large differences. In the most pessimistic views, a few locations enjoy the benefits of the frontier. The role of public policy aims at reaching a societal

ideal—orienting toward ameliorating unequal economic outcomes caused by unequal supply.

8.4.1 Informing Open Questions about Subsidies

Differences in these two views animate many open research questions today about optimal subsidies for digital infrastructure. Examples discussed in this chapter illustrate why debates examine the same fact base but point in different directions. The experience with CDNs supports the optimistic view, because of the wide geographic dispersion of supply and the presence of third parties. In contrast, the experience with data centers supports the less optimistic view, because of the concentration of supply around urban cities and the persistent demand for local supply. The two views also differ in their interpretation of the diffusion of broadband, with one side stressing the speed with which it reached a high percentage of households, and the other lamenting the slowing rate of adoption. These examples suggest no general answer will emerge, because analysis depends on specific cost conditions and use cases, and these are moving targets.

Deeper research can address some of the tension. For example, though users prefer a local supply of infrastructure when it is available, it may be possible to use remote data centers, cloud storage, and/or satellites. A similar trade-off faces users choosing between satellites and wireline broadband. These topics would be, and could be, informed by estimates of demand.

Another challenge arises from government efforts to employ infrastructure for noneconomic outcomes, such as informing citizens, furthering the education of children, contributing to the public health of a local population, or guaranteeing its safety. How much does society want to spend to encourage those goals? How much should it pay to build out the internet in low-density places to organizations with public missions, such as libraries, schools, hospitals, and public dispatchers? It can be challenging to translate these demands into pecuniary terms. As with the demand for many public goods, it is naive to presume an easy answer. With a moving frontier of acceptable quality, the measurement issues are especially vexing.

These challenges animate debates about government support for digital infrastructures, which tend to divide into one of three categories. The first category has to do with services at schools, hospitals, libraries, or governmental organizations such as police, fire, and other government services. The second concerns services for business. The third focuses on households.

The first category contains the most difficult issues to measure but, ironically, has tended to contain the least vociferous debate. Political systems in developed countries tend to view these investments as urgent. It is common for the entities in a locality to coordinate their purchases. Considerable federal funding has been redirected from universal services funds toward these use cases. That said, there is a largely open research question related to the

economic benefits of such investment. Analysis similar to Athey and Stern (2002) is exemplary, and there could be much more.

The second category, subsidies to business, also contains difficult measurement issues because of externalities. The effects from building new broadband, for example, may shape wages and employment of the local populace, and the building may shape the accumulation of additional services built on top of digital infrastructure. These debates are particularly fraught because of issues inferring causation—that is, does cheaper or better broadband *cause* better economic outcomes? The answers to those questions depend on counterfactual analysis. On the one hand, how would business behave in its procurement of services in the absence of subsidy? If business would have paid for broadband, why should the government subsidize it? On the other hand, what if the purchases take place in a location that faces eroding economic growth prospects? Would more, cheaper, or better broadband slow the erosion of the underlying economic value of economic activities performed in nearby areas or prevent an economic decline in the absence of such subsidy?

Such counterfactual questions frame difficult challenges for policy assessment. In the context of a backward-looking assessment of subsidies to broadband build-out in low-density settings, the analyst asks, What would have happened to wages, employment, and other indicators of economic prosperity in the absence of subsidy? An ideal empirical experiment would compare two otherwise similar locations which differ in only one respect: one received a subsidy and the other did not. Historical circumstances rarely produce such comparisons, however, so research has to find clever ways to exploit the minutiae of such situations, as in Boik (2017) or Skiti (2019).

Forward-looking policy assessment requires even more information. Analysis needs to understand the accumulation of business activities in the same location, which are more likely to generate large regional gain. For example, many localities and state development agencies grant tax abatements to build data centers, hoping to foster more development in a locale. While that might generate construction jobs for a short period, it is less obvious that such structures make a location inviting for further development on the digital frontier. Even a moderately large data center does not employ many people. What are the spillovers for the local service economy? What is the evidence about accumulation and spillovers after such tax abatements? These are open questions.

The third category, subsidies for access to households in low-density settlements, tends to focus on settings in which the costs of access are high and incremental gains go to a small number of households. This situation arises often in areas experiencing spotty housing settlement, such as rural areas. What would the households do with or without subsidy? What are the incremental gains in one situation compared with the other? Once again,

the analysis of gains depends on counterfactual questions, and these are not easily answered.

All these questions lack general answers because analysis depends on achievable aims, which vary over time. The costs and capabilities for satellite service, for example, have changed considerably over the past two decades, and so too have the capabilities of both fixed-wireless and mobile wireless communications over long distances. Economic circumstances and prospects also vary considerably across the thousands of low-density counties of the US, so the marginal potential adopter can and will vary for any specific proposal, and so too will the cost-benefit calculation.

In developing countries, all three categories of questions arise in different forms, once again preventing the emergence of any durable general answer. For example, a developing country may consider subsidies for capital expenditures for wireless infrastructure, which provides a foundation for further building of public services, business, and household activities. The counterfactual questions are especially challenging because investments in infrastructure alone may be insufficient to generate economic growth—that is, in the absence of complementary physical and human capital, particularly digital infrastructure and related firm investments in commercialization (Cirera and Maloney 2019). At the same time, the presence of network effects in wireless devices (Björkegren 2019) implies that subsidies might be beneficial well beyond the gains affiliated with satisfying the initial demand. Building wireless use could jump-start the use of digital infrastructure, which can become the basis for the development of further applications, such as in basic household finance and microloans for small entrepreneurs. Once again, the open question is fundamental: What is the evidence about accumulation and spillovers after building wireless digital infrastructure?

8.4.2 Open Questions about Governance

Historical studies of success and failure to innovate could illuminate understanding about how governance shapes the evolution of digital infrastructure. For example, the supply of infrastructure has supported the growth of valuable, standardized, and large-scale communications, such as texting, email, and a browser-supported advertising-oriented media market. That growth has fostered availability on multiple devices, supported by both wireline and wireless access. Any supplier can find the relevant technical standards and build a component that interoperates with the existing system. Why did that emerge, and what pitfalls were avoided? Similar questions arise about privacy standards. Why does this system work well for some attributes and not others?

The examples presented here will suggest that the boundary between public and private is in flux across a wide set of activities. Some software is private, some is open-source, and some employs a mixed model. Some software comes from consortia, other software from standard-setting orga-

nizations, and still other software from private suppliers. Government policy plays a variety of roles—for example, in subsidizing research and invention, in workforce training in higher education, in providing some services, and in defining legal boundaries for different types of organizations. The precise boundaries are open to debate and differ substantially across countries. Do those boundaries matter for economic outcomes? Answering that question could inform policy debates in many countries.

The governance of ubiquitous software requires attention, though such activity falls far outside the scope of (traditionally) regulated markets or public goods. Concrete examples can illustrate the open questions. Contrast two starkly different models. Some privately supplied software has achieved ubiquitous use, such as the Microsoft Operating System, Oracle Server, and Android/iPhone smartphone operating systems. Private firms supply this software, upgrade it, service requests, and exclude those who fail to pay an appropriate price. Another model also yields ubiquitous software. The World Wide Web Consortium is one example. Managed by a not-for-profit consortium, which regularly upgrades the software, the Web achieves ubiquity through nonexclusion, making upgrades available without restriction. The continuing success of the Web illustrates a model that leads to widespread use. What economic factors lead to a match between these governance models and market settings? How much difference does the governance model make to outcomes?

Web server software raises similar questions and offers different insights. Different users today largely employ three different servers: Apache, IIS, and Nginx. The first one descended from earliest experiments with web servers at the University of Illinois, organized as an open-source project. Microsoft offers IIS, generally as part of a range of the enterprise software it offers and certifies. The third, Nginx, comes in a freemium form, with a fast but limited version available at no charge. An enterprise version requires payment for services. Apache and IIS had a large impact on the market in the first two decades, but Nginx has enabled large gains in high-volume servers. The trade-offs between each of these organizational forms defy easy characterization.

As an example, Athey and Stern (2015) ask why some countries use more pirated operating system software. Their framework contrasts two broad determinants: (1) variation in willingness to pay for software, which shapes economic incentives to pirate software, and (2) institutional enforcement of property rights, which shaped incentives for private actors to invest in software. Athey and Stern measure the former with economic variables, such as per capita income, while they measure the latter with country-specific histories of respecting property rights. If the former is important, then sellers of proprietary software could potentially change their pricing strategies. If the latter is important, then pricing is unlikely to address the challenge, and better enforcement of legal regimes for property rights could have a larger effect.

Their framework provides a pathway forward. Two differences between operating systems and other internet infrastructure potentially shape the economics of other digital infrastructure. In most settings, infrastructure must be available for continuous operations and compatible with other parts of the internet. Continuous operation and compatibility requires the range of complementary operations mentioned previously.

The definition of infrastructure remains fluid in widely used software tools as well. For example, consider software repositories, such as GitHub, which has become common. GitHub aids the sharing of code and reduced frictions in large-scale projects. Making collaboration across distance easier, GitHub's creation had a well-known productivity impact; and Microsoft recently purchased the entire platform for \$7.5 billion in stock. It is essential infrastructure for many software projects. How Microsoft's purchase shapes GitHub's productive impact remains an open question.

Mapping software offers an example of the new frontier of the public-private boundary. The fundamental work was once thought to be solely a government function, and private firms merely repackaged the information in more accessible format for general consumption. At present, however, digital mapping has passed to either proprietary or open-source projects, which draw input from crowds, and these compete with one another. These platforms vary in their governance and source of input, as well as in response to new opportunities (see, e.g., Nagaraj 2019; Nagaraj and Piezunka 2018). The next generation of mapping for autonomous vehicles has moved to private sources. Different firms use distinct models of how to use input from users. All mapping depends critically on government-funded satellites that provide GPS (Global Positioning System) coordinates, so public support is never far away. Are the incentives to develop digital maps too low or too high, and do they result in too few or too many development projects? What are the incentives to share results, once they are developed?

A different and important insight comes from research focused on now-casting in developing countries—that is, using present economic activity to forecast events over a short time horizon, particularly where GDP measurement apparatus is absent or primitive. Near ubiquitous digital infrastructure can offer a way forward in measurement. For example, Indaco (2020) uses Twitter activity (as measured through GPS-labeled photos) to determine whether geo-located IP addresses give as much information as light from satellite photos. The study correlates Twitter use with other measures of economic activity, such as the light from satellite photos, because the same types of advanced investments support both—namely, continuous electrical supply, skilled labor, and a range of complementary investments. Ackermann and Angus (2014) provide a similar exercise when they examine the distribution of IP addresses. Once again, this provides evidence of economic activity.

To close, consider this provocative question: Is Wikipedia digital infra-

structure? Its ubiquity suggests it ought to be treated as such. It receives more than 15 billion pages views per month.⁴⁶ At the time of this writing, over 5.9 million articles grace its web pages in English alone, with more than 500 new articles added each day. Volunteers built the entire corpus of text. More to the point, because of its not-for-profit status, aspiration toward a neutral point of view, and minuscule storage and transmission costs, the scale economies appear virtually limitless. Wikipedia has become a focal site on which many others depend, including many search engines and Q&A sites. Many software firms also use it to complement their documentation efforts on GitHub, providing longer explanations and links.

The Wikipedia example epitomizes the open questions of this topic: What is and is not infrastructure when public funding is absent? Where are the boundaries of public and private when the private infrastructure contains properties similar to public goods? Can something be called infrastructure merely if it is shared, inexpensive, nonexclusive, and seemingly essential? Is the source—either public or private—relevant to the economics or virtually irrelevant?

To finish, note that Wikipedia remains unavailable in China, where the government firewall blocks access. That is but one example of many that illustrates the “splintering” of the internet. That occurs as a result of the erosion of compatibility of complementary equipment and software, resulting in distinct regions of the globe pursuing their own direction of technical developments, each internally consistent within national boundaries, yet inconsistent and incompatible across borders. Splintering has begun to arise as different governments censor content and impose limits on the operations of applications consistent with local preferences for privacy, security, copyright, and other government policy. Some of these actions have begun to migrate into the infrastructure layers, where governments impose, for example, distinctly different packet-inspection processes in routers or different back-door designs within operating systems. These actions and policies frame open questions about the consequence for seamless interoperability.

8.5 Conclusion

Long before it spread across the globe, it was fashionable to call the internet an “information superhighway.” The label arose from a combination of observation and aspiration. The *observation* contained a grain of truth about the physical layout of the internet. Many backbone lines followed existing rights-of-way for roads, bridges, and highways. The *aspiration* channeled a proposed vision for the future, with the government subsidizing the capital expenditure and leaving the assets unpriced—as with a freeway. The

46. “Wikimedia Statistics,” <https://stats.wikimedia.org/v2/#/all-projects>.

aspiration advanced an ideal in which information remained unpriced and subsidized by government support.⁴⁷

With the benefit of decades of hindsight, we can see that both the observation and aspiration about highways bear only partial resemblance to the present state of commercial digital infrastructure. One can gain insight from understanding the merits and shortfalls of the comparison.

Begin with the similarities. In the past few decades economic actors shared the use of the long-lasting capital of digital infrastructure, like much other infrastructure, and many economic actors employed digital infrastructure as an intermediate input in the production of goods and services. As an intermediate input, digital infrastructure acted much like a new road connecting two areas with previously poor connections, lowering frictions between potential transactions in different locations (Goldfarb and Tucker 2019). At a high level of abstraction, lowering of frictions created two types of new economic opportunities, either fostering cooperative agreement between suppliers of complementary inputs or encouraging competition from suppliers who serve new customers in new areas. New supply chains and applications built on these, such as those that employed more personalization; these would not have emerged or deployed in the absence of the low-cost and reliable network infrastructure.

The metaphor goes only so far in illuminating the core of the economic challenges, however. The building of digital infrastructure was not a one-time event, and the continued improvement changed transactions along many dimensions. The volume and type of traffic grew, and the degree of personalization increased, which, in turn, changed the viability of different services and the prospects for the firms offering services. The contrast with roads and highways could not be sharper: Roads typically do not undergo improvements in their key attributes every few years, while in digital infrastructure innovations accumulated, from many different suppliers, producing a system with capabilities that no central planner or brilliant designer could have specified in advance. New digital infrastructure supported frequent reassessments; in turn, these encouraged more experiments to expand service. That pace of improvement also heightened disagreements among distinct views about how to make valuable use of opportunities enabled by improved infrastructure. That enabled “innovation from the edges” (Greenstein 2015), which raised the importance of bringing frontier applications to market to settle the unresolved question about how to create value.

More pointedly, digital infrastructure *does not* resemble roads and highways in its pricing or governance. Building and operating roads and highways are largely government functions, and, relatedly, most highways and surface

47. The aspiration became associated with a range of policy initiatives that subsidized the internet for research in the late 1980s and, eventually, with the specific aspirations of a presidential candidate, Al Gore. See, for example, Greenstein (2015), chaps. 2 and 3, for a discussion of these policies.

streets remain unpriced and nonexcluded, with the possible exception of some toll roads and bridges. Questions about design and enhancement become decisions in the public sector. In contrast, while the role of public funding was once important, today private funding lags behind investment in the vast majority of digital infrastructure. Public funding continues to play a role in R&D activities, and the economic justifications for those subsidies are strong as a result of the externalities for suppliers and users. Beyond that, however, the degree of government intervention in digital infrastructure differs substantially from the typical practices with roads. The level of tax subsidy is much lower for digital infrastructure and more haphazard, affiliated with local tax abatement for large projects, zoning for new access services, or subsidies for internet access in costly areas. Modern suppliers face minimal mandates to become ubiquitous, reliable, and inexpensive beyond what market forces incentivize them to build and perform, whether governments accommodate those incentives or not. Despite the societal importance of fostering widespread use of frontier services, providers have unfettered discretion over price and other aspects of service. Failure by users to pay the minimal price leads to denying service to users. Altogether that situates the boundary between public and private at a substantially different place.

As of this writing, many of the most fundamental economic questions still remain unanswered. As this chapter has stressed, variance in the supply and use of innovative digital infrastructure arises within every developed country, as well as between developed and developing countries. Much of that variance arises because of differences in commercial incentives—quite unusual for infrastructure with such recognized importance for economic outcomes. That sets the stage for numerous research questions about the rate and direction of those incentives, as well as whether infrastructure’s performance achieves societal goals.

References

- Abbate, Janet. 1999. *Inventing the Internet*. Cambridge, MA: MIT Press.
- Ackermann, Klaus, and S. D. Angus. 2014. “A Resource Efficient Big Data Analysis Method for the Social Sciences: The Case of Global IP Activity.” *Procedia Computer Science* 29: 2360–69.
- Ahlfeldt, Gabriel, Pantelis Koutroumpis, Tommaso Valletti. 2017. “Speed 2.0: Evaluating Access to Universal Digital Highways,” *Journal of the European Economic Association* 15 (3): 586–625.
- Anderson, Monica, Andrew Perrin, Jingjing Jiang, and Madhumitha Kumar. 2019. “10% of Americans Don’t Use the Internet. Who Are They?” Pew Research Center, April. <https://www.pewresearch.org/fact-tank/2019/04/22/some-americans-dont-use-the-internet-who-are-they/>.
- Athey, Susan, and Scott Stern. 2002. “The Impact of Information Technology on Emergency Health Care Outcomes.” *Rand Journal of Economics* 33 (3): 399–432.

- . 2015. “The Nature and Incidence of Software Piracy: Evidence from Windows.” *Economic Analysis of the Digital Economy*, edited by A. Goldfarb, S. Greenstein, and C. Tucker. Chicago: University of Chicago Press.
- Bar, Talia, and Aija Leiponen. 2014. “Committee Composition and Networking in Standard Setting: The Case of Wireless Telecommunications.” *Journal of Economics & Management Strategy* 23 (1): 1–23.
- Baron, Justus, Kirti Gupta, and Brandon Roberts. 2018. “Unpacking 3GPP Standards.” *Journal of Economics & Management Strategy* 27 (3): 433–61.
- Björkegren, Daniel. 2019. “The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda.” *Review of Economic Studies* 86 (3): 1033–60.
- Boik, Andre. 2017. “The Economics of Universal Service: An Analysis of Entry Subsidies for High Speed Broadband.” *Information Economics and Policy* 40: 13–20.
- Boik, Andre, Shane Greenstein, and Jeff Prince. 2019. “The Persistence of Broadband User Behavior: Implications for Universal Service and Competition Policy.” *Telecommunications Policy* 43 (8): 101820.
- Bottger, Timm, Felix Cuadrado, Gareth Tyson, Ignacio Castro, and Steve Uhlig. 2018. “Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN.” *ACM SIGCOMM Computer Communications Review* (CCR) 48 (1).
- Brynjolfsson, Erik, Avinash Collis, and Felix Eggers. 2019. “Using Massive Online Choice Experiments to Measure Changes in Well Being.” *PNAS* (Proceedings of the National Academy of Sciences) 116 (15): 7250–55.
- Brynjolfsson, Erik, and Joo Hee Oh. 2012. “The Attention Economy: Measuring the Value of Free Digital Services on the Internet.” 33rd International Conference on Information Systems.
- Burnham, Brad, Shane Greenstein, Neil Hunt, Kevin McElearney, Marc Morial, Dennis Roberson, and Charles Slocum. 2013. *Issues in Data Caps and Usage-Based Pricing*. Report, Federal Communications Commission, Open Internet Advisory Committee, Economic Impacts Working Group, May 5. <https://www.hbs.edu/faculty/Pages/item.aspx?num=49733>.
- Byrne, David, and Carol Corrado. 2019. “Accounting for Innovation in Consumer Digital Services: IT Still Matters.” NBER Working Paper 26010. Cambridge, MA: National Bureau of Economic Research.
- Byrne, David, Carol Corrado, and Dan Sichel. 2018. “The Rise of Cloud Computing: Minding Your P’s, Q’s and K’s.” NBER Working Paper 25188. Cambridge, MA: National Bureau of Economic Research.
- Cardona, M., T. Kretschmer, and T. Strobel. 2013. “ICT and Productivity: Conclusions from the Empirical Literature.” *Information Economics and Policy* 25 (3): 109–25.
- Cariolle, Joel. 2019. “Telecommunication Submarine-Cable Deployment and the Digital Divide in Sub-Saharan Africa.” CESifo Working Paper 7415. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3338769##.
- Cirera, Xavier, and William Maloney. 2019. *The Innovation Paradox: Developing-Country Capabilities and the Unrealized Promise of Technological Catch-Up*. Washington, DC: World Bank Group.
- Chen, Yongmin, and Scott Savage. 2011. “The Effects of Competition on the Price for Cable Modem Internet Access.” *Review of Economics and Statistics* 93 (1): 201–17.
- Clark, David. 2018. *Designing an Internet*. Cambridge, MA: MIT Press.
- Clark, David, Steven Bauer, William Lehr, kc claffy, Amogh Dhamdhere, Bradley

- Huffaker, and Matthew Luckie. 2014. "Measurement and Analysis of Internet Interconnection and Congestion." TPRC Conference Paper, September 9.
- Comin, Diego, and Bart Hobijn. 2010. "An Exploration of Technology Diffusion." *American Economic Review* 100 (December): 2031–59.
- Connolly, Michelle, and J. E. Priege. 2013. "A Basic Analysis of Entry and Exit in the US Broadband Market, 2005–2008." *Review of Network Economics* 12 (3): 229–70.
- Coyle, Diane, and David Nguyen. 2018. "Cloud Computing and National Accounting." ESCoE Discussion Paper, December. Office of National Statistics, UK.
- Cranor, Lorrie Faith, Serge Egelman, Steve Sheng, Aleecia M. McDonald, and Abdur Chowdhury. 2008. "P3P Deployment on Websites." *Electronic Commerce Research and Applications* 7 (3): 274–93.
- DeStefano, Tim, Richard Kneller, and Jonathan Timmins. 2018. "Broadband Infrastructure, ICT Use and Firm Performance: Evidence for UK Firms." *Journal of Economic Behavior & Organization*, 155(C): 110–39.
- . 2019. "Cloud Computing and Firm Growth." Working paper presented at NBER Winter Digitization Meetings.
- Edelman, Benjamin, and Michael Schwarz. 2015. "Pricing and Efficiency in the Market for IP Addresses." *American Economic Journal: Microeconomics* 7 (3): 1–23.
- Ewens, Michael, Ramana Nanda, and Matthew Rhodes-Kropf. 2019. "Cost of Experimentation and the Evolution of Venture Capital." Working Paper 15-070, Harvard Business School.
- Falck, Oliver, Robert Gold, and Stephan Heblich. 2014. "E-lection: Voting Behavior and the Internet." *American Economic Review* 104 (7): 2238–65.
- Fernandes, Ana, Aaditya Mattoo, Huy Nguyen, and Marc Shiffbauer. 2019. "The Internet and Chinese Experts in the Pre-Ali Baba Era." *Journal of Development Economics* 138: 57–76.
- Forman, Chris. 2005. "The Corporate Digital Divide: Determinants of Internet Adoption." *Management Science* 51 (4): 641–54.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2005. "How Did Location Affect Adoption of the Commercial Internet? Global Village vs. Urban Leadership." *Journal of Urban Economics* 58 (3): 389–420.
- . 2012. "The Internet and Local Wages: A Puzzle." *American Economic Review* 102 (1): 556–75.
- Goldberg, Samuel, Garrett Johnson, and Scott Shriver. 2019. "Regulating Privacy Online: The Early Impact of the GDPR on European Web Traffic and E-Commerce Outcomes." <https://ssrn.com/abstract=3421731>.
- Goldfarb, Avi, and Jeff Prince. 2008. "Internet Adoption and Usage Patterns Are Different: Implications for the Digital Divide." *Information Economics and Policy* 20 (1): 2–15.
- Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57 (1): 3–43.
- Goolsbee, Austan, and Peter Klenow. 2006. "Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet." *American Economic Review* 96 (2): 108–13.
- Greenstein, Shane. 2015. *How the Internet Became Commercial: Innovation, Privatization, and the Birth of a New Network*. Princeton, NJ: Princeton University Press.
- . 2020. "The Basic Economics of Internet Infrastructure." *Journal of Economic Perspectives* 34 (2): 192–214.
- Greenstein, Shane, and Ryan McDevitt. 2011. "The Broadband Bonus: Estimating Broadband Internet's Economic Value." *Telecommunications Policy* 35: 617–32.

- Greenstein, Shane, and F. Nagle. 2014. "Digital Dark Matter and the Economic Contribution of Apache." *Research Policy* 43: 623–31.
- Greenstein, Shane, and Tommy Pan Fang. 2019. "Where the Cloud Rests: the Economic Geography of Data Centers." Working paper, Harvard Business School.
- Greenstein, Shane, Martin Peitz, and Tommaso Valletti. 2016. "Net Neutrality: A Fast Lane to Understanding the Trade-Offs." *Journal of Economic Perspectives* 30 (2): 127–50.
- Haller, Stephanie, and Sean Lyons. 2015. "Broadband Adoption and Firm Productivity: Evidence from Irish Manufacturing Firms." *Telecommunications Policy*: 1–15.
- . 2019. "Effects of Broadband Availability on Total Factor Productivity in Service Sector Firms: Evidence from Ireland." *Telecommunications Policy* 43 (1): 11–22.
- Hitt, Lorin, and Prasanna Tambe. 2007. "Broadband Adoption and Content Consumption." *Information Economics and Policy* 19 (3–4): 362–78.
- Hjort, Jonas, and Jonas Poulsen. 2019. "The Arrival of Fast Internet and Employment in Africa." *American Economic Review* 109 (3): 1032–79.
- Horrigan, John. 2020. "Measuring the Gap: What Is the Right Approach to Exploring Why Some Americans Do Not Subscribe to Broadband?" National Data Inclusion Alliance. <https://www.digitalinclusion.org/measuring-the-gap/>.
- Huston, Geoff. 2017. "The Rise and the Rise of Content Distribution Networks." Presentation at Workshop on Information Economics, Center for Applied Internet Data Analysis, University of California, San Diego, December 2017. http://www.caida.org/workshops/wie/1712/slides/wie1712_ghuston.pdf.
- Indaco, Augustin. 2020. "From Twitter to GDP: Estimating Economic Activity from Social Media." *Regional Science and Urban Economics* 85 (November): 103591. <https://www.sciencedirect.com/science/article/pii/S0166046220302763>.
- Jin, Wang, and Kristina Steffensen McElheran. 2018. "Economies before Scale: Learning, Survival, and Performance of Young Plants in the Age of Cloud Computing." Rotman School of Management Working Paper 3112901. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3112901.
- Knieps, G., and J. M. Bauer. 2016. "The Industrial Organization of the Internet." In *Handbook on the Economics of the Internet*, edited by M. Latzer and J. M. Bauer, 23–54. Cheltenham, UK: Edward Elgar.
- Kolko, Jed. 2012. "Broadband and Local Growth." *Journal of Urban Growth* 71 (1): 100–113.
- Malone, Jacob, Aviv Nevo, and Jonathan W. Williams. 2017. "The Tragedy of the Last Mile: Economic Solutions to Congestion in Broadband Networks." NET Institute Working Paper 16-20. <https://ssrn.com/abstract=2849869>.
- McElheran, Kristina. 2015. "Do Market Leaders Lead in Business Process Innovation? The Case(s) of E-Business Adoption." *Management Science* 61 (6): 1197–216.
- McManus, Brian, Aviv Nevo, Zachary Nolan, and Jonathan W. Williams. 2018. "Steering Incentives and Bundling Practices in the Telecommunications Industry." NET Institute Working Paper 18-12. <https://ssrn.com/abstract=3267060>.
- Nagaraj, Abhishek. 2019. "Information Seeding and Knowledge Production in Online Communities: Evidence from OpenStreetMap." Paper, UC Berkeley Haas School of Business. http://abhishek.com/files/nagaraj_seeding.pdf.
- Nagaraj, Abhishek, and Henning Piezunka. 2018. "Deterring the New, Motivating the Established—The Divergent Effect of Platform Competition on Member Contributions in Digital Mapping Communities." http://abhishek.com/files/nagaraj_piezunka_competition.pdf.

- Nagle, Frank. 2019. "Government Technology Policy, Social Value, and National Competitiveness." Working Paper 19-103, Harvard Business School.
- National Telecommunications and Information Administration. 2011. "Digital Nation: Expanding Internet Usage." February 2011. https://www.ntia.doc.gov/files/ntia/publications/ntia_internet_use_report_february_2011.pdf.
- Nevo, Aviv, John Turner, and Jonathan Williams. 2016. "Usage-Based Pricing and Demand for Residential Broadband." *Econometrica* 84 (2): 411–43.
- Norton, William B. 2014. *The Internet Peering Playbook: Connecting to the Core of the Internet*. Palo Alto, CA: DrPeering Press.
- Nuechterlein, Jonathan, and Philip Weiser. 2005. *Digital Crossroads: Telecommunications Law and Economics in the Internet Age*. Cambridge, MA: MIT Press.
- OECD. 2014. *Measuring the Digital Economy: A New Perspective*. Paris: OECD Publishing. <https://doi.org/10.1787/9789264221796-en>.
- Poliquin, Chris. 2020. "The Wage and Inequality Impacts of Broadband Internet." Working paper, University of California, Los Angeles. https://poliquin.xyz/files/poliquin_jmp.pdf.
- Rosston, Gregory, Scott Savage, and Donald Waldman. 2010. "Household Demand for Broadband Internet Service in 2010." *B. E. Journal of Economic Analysis & Policy* 10 (1): 79.
- Rosston, Gregory L., and Scott J. Wallsten. 2020. "Increasing Low-Income Broadband Adoption through Private Incentives." Stanford Institute for Economic Policy Research Working Paper 20-001. <https://siepr.stanford.edu/research/publications/increasing-low-income-broadband-adoption-through-private-incentives>.
- Seamans, Robert. 2012. "Fighting City Hall: Entry Deterrence and Technology Deployment in the Cable TV Industry." *Management Science* 58 (3): 461–75.
- Simcoe, Timothy. 2012. "Standard Setting Committees: Consensus Governance for Shared Technology Platforms." *American Economic Review* 102 (1): 305–36.
- Skiti, Tedi. 2019. "Strategic Technology Adoption and Entry Deterrence in Broadband." *Industrial and Corporate Change* 29 (3): 713–29.
- Wallsten, Scott, and Colleen Mallahan. 2013. "Residential Broadband Competition in the United States." In *The Economics of Digitization*, edited by Avi Goldfarb, Shane Greenstein, and Catherine Tucker. Cheltenham, UK: Edward Elgar.
- Wallsten, Scott, and James Riso. 2010. "Residential and Business Broadband Prices, Part 2: International Comparisons." Policy paper, Technology Policy Institute. https://techpolicyinstitute.org/policy_paper/residential-and-business-broadband-prices-part-2-international-comparisons/.
- Wang, Shuang, Jacob LaRiviere, and Aadharsh Kannan. 2019. "Spatial Competition and Missing Data: An Application to Cloud Computing." Working paper presented at NBER Winter Digitization Meetings.
- Whitacre, Brian, Roberto Gallardo, and Sharon Strover. 2014. "Broadband's Contribution to Economic Growth in Rural Areas: Moving towards a Causal Relationship." *Telecommunications Policy* 38 (11): 1011–23.
- World Economic Forum. 2016. *The Global Information Technology Report*. <https://www.weforum.org/reports/the-global-information-technology-report-2016>.
- Zhuo, Ran, Bradley Huffaker, kc claffy, and Shane Greenstein. 2019. "The Impact of the General Data Protection Regulation on Internet Interconnection." NBER Working Paper 26481. Cambridge, MA: National Bureau of Economic Research.