Journal of Economic Perspectives—Volume 34, Number 2—Spring 2020—Pages 192–214

The Basic Economics of Internet Infrastructure

Shane Greenstein

This internet barely existed in a commercial sense 25 years ago. In the mid-1990s, when the data packets travelled to users over dial-up, the main internet traffic consisted of email, file transfer, and a few web applications. For such content, users typically could tolerate delays. Of course, the internet today is a vast and interconnected system of software applications and computing devices, which society uses to exchange information and services to support business, shopping, and leisure. Not only does data traffic for streaming, video, and gaming applications comprise the majority of traffic for internet service providers and reach users primarily through broadband lines, but typically those users would not tolerate delays in these applications (for usage statistics, see Nevo, Turner, and Williams 2016; McManus et al. 2018; Huston 2017). In recent years, the rise of smartphones and Wi-Fi access has supported growth of an enormous range of new businesses in the "sharing economy" (like, Uber, Lyft, and Airbnb), in mobile information services (like, social media, ticketing, and messaging), and in many other applications. More than 80 percent of US households own at least one smartphone, rising from virtually zero in 2007 (available at the Pew Research Center 2019 Mobile Fact Sheet). More than 86 percent of homes with access to broadband internet employ some form of Wi-Fi for accessing applications (Internet and Television Association 2018).

It seems likely that standard procedures for GDP accounting underestimate the output of the internet, including the output affiliated with "free" goods and the restructuring of economic activity wrought by changes in the composition of firms who use advertising (for discussion, see Nakamura, Samuels, and Soloveichik

■ *Shane Greenstein is the Martin Marshall Professor of Business Administration, Harvard Business School, Boston, Massachusetts. His email address is sgreenstein@hbs.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.34.2.192.

2016, or the Spring 2017 symposium in this journal with articles by Feldstein 2017; Syverson 2017; Groshen et al. 2017). To illustrate the magnitude of the measured economic changes, online advertising contributed $105.9 billion in revenue to the GDP in the categories of Internet Publishing and Broadcasting as well as Web Search Portals in 2017, which had grown 250 percent in the previous five years. The Census Bureau estimates electronic retailing at over $545 billion for just electronic shopping and mail order houses (NAICS 4541), a growth of 65 percent over the same period (based on the Census data on Statistics of US Business).

The external face of the internet has become part of everyday life. However, the internal structure and operation of the internet have remained largely invisible, both to the public and to most economists. This essay will begin by discussing the processes that support delivery of internet services. The internet's infrastructure contains many different types of equipment: root servers, fiber, broadband lines, networking switches and routers, content delivery networks, cellular towers, and others. Meanwhile, the internet's "backbone" consists of enormous data lines, specifically, the lines that interconnect networks and core routers for transmitting packets of data. Other elements of the internet infrastructure include cloud facilities and the parts of the internet that have been taken inside large firms like Google and Amazon. With an understanding of the mechanics of the internet structure, it becomes possible to address questions like: What determines the pricing and terms for exchanging data? What determines the incentives for improving infrastructure? How evenly spread is frontier digital infrastructure across regions?

The discussion will illustrate some classic issues in the economics of networks. Networks which have an agreed upon set of standards and rules can be self-perpetuating in a wide range of circumstances because existing users and potential new ones will be attracted to the well-established network. However, when a network involves both multiple end-users and multiple players within the network who can impose costs and fees on each other, there may be times when negotiations may threaten to deadlock. Expanding a network in its original form may be fairly straightforward, but more complex changes to the operation of a network can be problematic, both because such changes may threaten to disturb the shared rules that make the network function and because the players who would need to invest in the change may find that they are not able to recoup a sufficient share of the benefits from other players in the market to make it worthwhile.

The discussion will focus on practices in a North American context and will oversimplify the explanations of its engineering. However, it should generate an understanding of how internet infrastructure works as well as it does. Technical terms will be introduced and explained as they arise. Additionally, Table 1 provides a glossary.

How Does Internet Data Travel?

To understand how the internet connects so many devices, let's start with a basic example: how a single user request for information from, say, Wikipedia, generates

Table 1
Glossary of Some Internet Terminology

Term	*Definition*
Backbone	The long distance and high capacity routes between interconnected networks and core routers in the internet.
BGP	Border Gateway Protocol. The most commonly used protocol for routing traffic on the internet and is one among many governing how network switches and servers send packets of data through the network. The most recent draft dates to 2006. See https://tools.ietf.org/html/rfc4271.
Broadband	Any high-speed internet access that is always on and faster than dial-up access.
CDN	Content Distribution Network. A distributed system of computers that acts as an intermediary for original content, and delivers content transparently to end users.
Cloud computing	An evolving model for enabling a ubiquitous and on-demand shared pool of configurable computing resources. Users typically provision these quickly.
Collocation facility	A location in which servers and other computing hardware reside.
DNS	Domain Name System is set of naming and numbering rules for affiliating common words with IP addresses, consistent with TCP/IP. Today ICANN oversees the system used on the internet.
DOCSIS	Data Over Cable Service Interface Specification. Developed by Cable Labs for cable system delivery of internet access (see discussion in Knieps and Bauer 2016, and Clark 2018).
DSL	Digital Subscriber Line. A form of broadband access retrofitted to telephone lines.
ICANN	The Internet Corporation for Assigned Names and Numbers. The nonprofit organization responsible for coordinating the maintenance and procedures of several databases related to the namespaces and numerical spaces of the internet. https://www.icann.org/.
IEEE	The Institute of Electrical and Electronic Engineers is a global association and organization of professionals working toward the development, implementation, and maintenance of technology-centered standardized products and services. https://www.ieee.org/.
IP address	Internet Protocol. Every device on the internet must have this numerical label assigned to it.
IXP	Internet Exchange Point. Typically a building operated by carriers or by a third party and configured for carrier colocation and interconnection of data traffic.
Protocol stack	The software that implements a family of protocols. These define a set of rules and regulations that determine how data transmits in telecommunications and computer networking.
TCP/IP	Transmission Control Protocol/Internet Protocol. This packet-switching protocol defines how to assemble packets, defines addresses when networks connect to each other, and is a family of protocols that determines the format and error correction processes for packets of data in the internet.
Wi-Fi	It is *not* wireless-fidelity, but is a set of protocols used by wireless routers and based around an IEEE 802.11 family of standards.

a number of instantaneous actions. This involves an explanation of the mechanics of moving data, typically unseen by the user.

Figure 1

Average Weighted Load Time Compared with Advertised Download Speed Federal Communications Commission (December 2018)

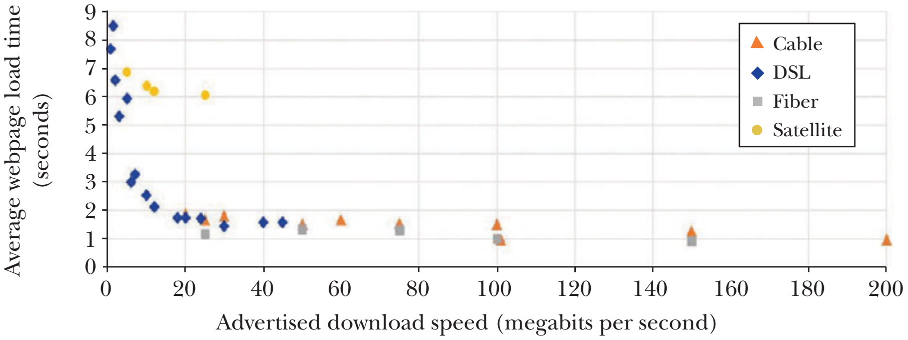

Source: Federal Communications Commission (2018).

Here is a simplified explanation of the mechanics: The user employs a web browser that has been installed on a computer, smartphone, or other web-enabled device. The user has access to an Internet Service Provider, or ISP. ISPs provide wireline or wireless access by building and operating the physical equipment that carries data from one place to another. The internet service provider takes the user's request to a name server. The name-server associates an internet protocol (IP) address with the requested destination—in this example, Wikipedia.org. Thus informed, the user's browser directs the query to the server with that IP address. Wikipedia's server responds by releasing the requested data in packets, which are formatted to comply with a specific protocol used to interconnect devices on the internet. That data travels to the user's ISP, which delivers it to the user's device, where it is rendered by the device into in a form the user can view.

Several different market transactions support this two-way flow of information. First, market transactions visibly determine the behavior of internet service providers who are typically paid on a monthly basis. There are broadly two types of ISPs: wireline and wireless providers. Wireline providers vary in their technology—listed here in order from slowest to fastest: Satellite, Digital Subscriber Line (DSL), cable modem, and fiber. Satellite in geostationary orbits deliver and receive data to and from almost any earthly location fitted with a "dish," which communicates with the satellite. DSL service is a retrofit on top of telephone lines to suit it to carrying data. Cable modem service involves the addition of switches and modems consistent with Data Over Cable Service Interface Specification (DOCSIS), which adds data services to cable television systems. Fiber typically involves newly laid lines of fiber optic wire to the customer.

Figure 1 shows a standardized test conducted by the Federal Communications Commission of several advertised tiers of speeds from 17 companies whose service is representative of the experience of the vast majority of US users. The data rates are expressed in megabits per second and translate speeds into a standardized user experience downloading web pages. As illustrated by Figure 1, users experience different download speeds from different tiers of advertised speeds for access technologies. Monthly prices vary accordingly. Typical satellite services cost $90 to $120 a month, on top of set up costs of at least $300 to $500. For DSL, monthly prices tend to range from $30 to $50 a month for only internet service. The largest provider of DSL services in the United States is AT&T, with approximately 16 million subscribers. Prices for cable modem service range from $50 to $80, depending on speed and data caps. The largest providers of cable modem is Comcast, with over 28 million customers. Prices for fiber to the home tend to range between $40 and $80 per month for only internet, depending on speed and data caps. The largest provider of fiber to premises and homes is Verizon Fios, with approximately 7 million subscribers. In any given location the set of options may be more limited to zero, one, or two wireline providers, plus a potential over-builder.

Wireless options differ in use from wireline broadband. While satellite service is available anywhere, most of its users are in low-density locations lacking wireline providers due to its low expense and speed. Estimates put the number of users at more than 8 million households in the United States. The largest providers of wireless services are the carriers Verizon Wireless and AT&T Wireless, with more than 150 and 160 million subscribers, respectively.

Another set of market transaction is invisible to users. The name-server firms are paid by the owners of domain names or their surrogate parties acting in the interests of website owners. The largest US name servers are Cloudflare, Amazon Web Services, and Akamai. While a name server may be a stand-alone firm, it has become increasingly common to offer name service bundled with other services, such as security. In addition, some organizations that send out large volumes of messages will provide their own in-house name server rather than paying for third-party services (for an explanation of this choice, see Bates et al. 2018).

Five Options for Data to Travel

With internet service providers and name servers playing their roles, one crucial step remains: how does the data actually travel between the internet service providers of the user and a content provider like Wikipedia and vice versa? Internet data can follow a multiplicity of paths between two points, which gives the system immense flexibility. How is the route for any given message determined? All options make use of the same routing tables and software protocols, which typically direct the packets of data to the least congested route. That process involves what is largely an engineering decision about how all networked participants collectively must behave in the presence of congestion on some routes. However, we will defer an explanation of how the prices for sending data are determined until later in the paper—because understanding the economics is more easily done after an explanation of the network's mechanics. For now, we will focus on the path taken by data as it goes from user to Wikipedia and back again. The data can travel between user and content provider by one of five options.

The first option is the simplest. If the user and the server contract with the same internet service provider, such as Comcast, then the data can be requested and delivered within the network of a single ISP. This path is common for bilateral communications between individuals, such as electronic mail—the majority of which involves two closely located participants. However, most other traffic, particularly traffic to support web and streaming applications, tends to involve content providers and recipients in far apart locations. Those interactions do not tend to stay within a single network due to the geographically fragmented and unconcentrated provision of US internet services providers.

That brings us to the most common current option in which the internet minimizes delays by rerouting a user request from server to content delivery networks (CDNs), which are geographically distributed networks of servers located near end users. Sometimes this is called "moving data to the edge of the network." Because such networks are physically close to users, CDNs reduce the response time. Many content providers choose to cache content at the CDN and update only the most timely and popular content so that most users are, in effect, exchanging content with the CDN rather than the ultimate provider of content. CDNs also can provide a layer of reliability and security: for example, even when some servers have gone down, the cached content in a CDN may keep a firm's content available for users. CDNs can also buffer content from a "denial-of-service" attack (in which an attacker seeks to disable a target by flooding it with messages).[1]

Content delivery networks did not exist at the outset of the commercial internet, but today, almost every commercial participant of any size employs them in some way for popular content. The largest third-party provider of CDN services in the United States is Akamai, with revenues of $2.7 billion in 2018. The next-largest provider of such services, Cloudflare and Limelight, had revenues of $192 million and $184 million in 2018, respectively. Though content delivery networks are unseen to the user, the vast majority of data received by a user comes directly via this route.

Three other options for moving data have been used for over two decades since the privatization of the internet (Greenstein 2015), but it is difficult to derive estimates of their frequency of use. In the distant past, all were more commonly used to move data from content providers directly to users—that is, without the intervention of a content delivery network. Today, these same forms move data from content providers to CDNs, complementing the CDNs in the vast majority of requests. These three forms—private peering, internet exchange points, and transit carriers—act as a substitute for the CDN in a smaller set of cases, as when the user requests unpopular content, or the content provider does not make an arrangement to employ a CDN.

---

[1] Readers may be interested in Patent 8613089B1, *Identifying a Denial of Service Attack in a cloud-based proxy service,* assigned to Cloudfare at https://patentimages.storage.googleapis.com/a0/90/f7/3f8aa8ef076cf4/US8613089.pdf.

"Private peering" arises when Wikipedia and the user (and the CDN supporting the user), have different internet service providers, but those two providers have a direct point of contact, and made a bilateral contract with each other to govern the exchange of data. In a typical contract, no money changes hands if over a month their data flows back and forth in rough proportion to each other. If one party gives a higher proportion of data to the other, the carrier who gives more data (on net) pays the other carrier for taking the traffic. Typically, these payments arise when traffic exceeds a negotiated ratio between four-to-one or eight-to-one—but no simple sentence can describe these contracts and negotiations (for discussion, see Norton 2014).

Two or more internet service providers also may exchange data at an internet exchange point (IXP), which may be run by a separate organization and configured as a place for carriers to meet and interconnect so they can exchange traffic. Each carrier pays a fee to the organization that houses the equipment that facilitates exchanging the data and may make numerous investments in the structures, backup energy, and equipment to keep these operating under all circumstances. Unlike private peering, all participants generally agree to send and take whatever volume of data their connection's capacity can handle. Charges may vary for each tenant and often has no relationship with volume of traffic. There are hundreds of IXPs in the United States and more all over the globe. The largest operator is Equinix, with over $5 billion in revenue and over 200 data centers in many cities, with some of these configured to serve as IXPs.

When the internet service providers for the user and for Wikipedia in our base example do not have any direct contact with each other, not even via an internet exchange portal, then a last possible form of making contact arises. One or more networks' lines acts as a transit carrier between the two ISPs. The carriers providing transit may have received compensation for that action depending on all their contracts with other carriers.

Incentives for Investment, Expansion, and Improvement

Notice an economic implication of this system: carriers have incentives to build more lines, make more connections, and relieve congestion, if and when it helps the firm to gain revenue or to avoid charges from other firms. Internet service providers face additional incentives to increase capacity and make connections if it enables them to increase revenue from users and/or avoid operational costs. These incentives appear to be consistent with a desirable long-term outcome— namely, more efficient and better options for routes to send and receive data. An interesting open question concerns the size of the private incentives in relation to the gains to the network. Transit lines are one component in a system, and improvements in one component confers benefits to all the other complementary components. Do most of the gains from better transit lines go to the content providers who use them, to the users who enjoy previously slower content, or to the internet service providers who may gain revenue from users for better services? The answer partly depends on pricing, which we discuss later.

A related question arises about the incentives to install content delivery networks. A third-party commercial content delivery network negotiates interconnection with an internet service provider or wireless access provider for the right to "collocate" a server close to users. The ISP or another network provider also may charge a "transit" fee to the CDN to take data over its network lines (that is, from the content firm's servers to the equipment installed by the CDN). The original content providers pay the CDN provider to redistribute content to users from the CDN's servers, which the content provider updates at an arranged schedule over the course of the day. This contractual arrangement arises in virtually any, albeit the smallest, ISPs in the United States, which suggests it serves the interest of ISPs.

Some large content providers, such as Google, Apple, Microsoft, Facebook, Amazon, and Netflix, operate their own content delivery networks and tailor the technical features to their own applications and services. Again, they negotiate a price that they pay to internet service providers for "collocation," and they sometimes pay fees for data transit. In practice, only large firms opt for this action because it is usually less expensive to contract with a third-party CDN for small to medium volumes of traffic. Also, for a number of reasons—scaling issues, negotiating frictions, and the collocation expense—some firms prefer to locate some of their private CDNs at internet exchange points, not within internet service providers.

It is an open question: Do most of the gains from better content delivery networks go to CDN providers who operate the servers, to the content providers who use them, to the users who enjoy previously unobtainable content, or to the internet service providers who charge collocation fees and also may gain revenue from users for better services? As with any network component, it is unclear how the private incentives compare with network-wide gains.

This question is important because the rise of content delivery networks was both a cause and symptom of changing user needs and dramatic network improvements. Many users have migrated to broadband with higher bandwidth, which increases user speeds. These users are more likely to desire and support new applications, which would have been infeasible without CDNs, such as "over-the-top" streaming services such as Netflix, Sling, Disney+, or HBO Go—that is, services that bypass cable or satellite television content and instead are provided directly to consumers over the internet.

The dramatic improvements are most visible in the heavy evolution of applications of the internet and the traffic that accompanies them. In the earliest days of the internet, text dominated traffic either in the form of email or passive browsing. By modern standards, the volumes of data were small in either direction. In contrast, households today receive increasingly many more magnitudes of data than they send, as the majority of traffic that households receive changes from static content to video and streaming (Huston 2017). For example, back in 2013, a median household used 20–60 gigabytes of data per month (Federal Communications Commission 2013). However, streaming a standard or high definition movie generates between 1 and 3 gigabytes per hour, far more data than any passive web-browsing ever could generate.

Merely binge-watching a single streamed series could massively increase household data use. Meanwhile, the largest streaming service, Netflix, has increased its US subscribership from 20 to 60 million over the second decade of the century, and it is far from the only streaming service. In short, as streaming of television and movies rises in households, the capacity of the underlying infrastructure to handle data-intensive applications must increase. It is always hard to answer the question of whether incentives to invest are optimal, but the experience of the internet certainly suggests that private incentives to invest have been sufficient to produce a dramatic expansion and upgrading of network components.

Data Centers and the Cloud

At the outset of the commercial internet, virtually all firms housed their servers on company premises. Businesses, however, eventually learned to gain scale economies by consolidating computing resources in one location, which gave birth to the data center. These structures contain many rows of servers on racks, matched to routine operations for support and maintenance of the internet. Designers eventually learned to configure these structures to house massive numbers of servers devoted to storage or computation, using architectural features that encourage low energy use and ensure reliable operations in the event of emergencies, among many features.

Some of the different ways that the market can send and receive information, such as peering and interconnection, also occur at some data centers configured for such a purpose. The inside wiring of a data center may support a specific set of activities. The data center for the New York Stock Exchange, for example, is located in New Jersey, and it 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—that is, 99.99 percent uptime—especially in critical functions that support transactions with sensitive customer data. These data centers may contain expensive backup generators, expensive structures to prevent flooding, and reinforcements in the floors to reduce any vibrations from passing vehicles. These expensive features pay off in certain situations; for example, due to 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.

Small data centers house tens of thousands of servers and can cost more than $100 million to build from scratch, while large data centers house hundreds of thousands of servers and can cost several billion dollars to build from scratch. One of the largest third-party facilities in the United States, the Lakeside Technology Center, resides two miles south of downtown Chicago in 1.1 million square feet of a converted building that formerly housed R.R. Donnelly's printing facilities. It is owned by Digital Realty Trust, a holding company that manages more than 200 data centers around the globe, generating just over $3 billion in revenue in 2018.

This building is an exception to the norm for data centers, which are typically large one-story buildings built on an expanse of land near abundant, inexpensive electricity and high-quality interconnection with the internet, often at a suburban location not far from the business users. 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.

Contracts for data centers cover every conceivable arrangement between ownership and rental markets. At one extreme, many buyers with generic needs—such as storage for backup—rent data center space, own the servers, and let others manage the building. At the other extreme, firms with unique computing needs—such as Facebook, Apple, Microsoft, Amazon, and Google—own and operate large private data centers and configure the building and servers to suit their applications.

A "cloud" service involves a data center that rents its services for storage, computing, or their respective applications to a service such as database, with the additional feature that users can turn the service turn off and on at will. The major cloud providers also increasingly offer additional software services for a nominal charge or none at all. For example, Amazon Web Services offers scores of cloud software services. Microsoft Azure supports many Microsoft products, such as Outlook, as a cloud service. Google offers Tensor Flow, a standard tool for machine learning, at no charge with its cloud service.

The demand for cloud services has grown as the services have improved and declined in price. Byrne, Carrado, and Sichel (2018) estimate a quality-adjusted price decline between 2009 and 2016 at 17.3 percent per annum for Amazon Web Services. Estimates of the growth in expenditure and market share within the industry depend on the precise definition of sales (for discussion, see Byrne, Carrado, and Sichel 2018; Coyle and Nguyen 2018), but some of the three biggest players are those just mentioned: Amazon Web Services (AWS), Microsoft Azure, and Google Cloud. In 2019, for example, AWS was widely regarded as the largest of these three cloud providers and brought in $35 billion of revenue—an increase of 40 percent from the prior year. The others are also growing rapidly. The appeal of cloud facilities comes from their flexibility, wide range of tools, and the option to substitute variable costs for fixed ones (Wang and McElheran 2017). It has enabled experimentation by many entrepreneurial applications (Ewens, Nanda, and Rhodes-Kropf 2019).

The private cloud providers increasingly use complex architectures to balance the loads from user demands—for example, using a mix of data centers for high-scale tasks and content delivery networks for rapid response for timely content. Cloud facilities provide updates to the CDNs at timely intervals and secondary response of less popular content, while home servers provide updates at slower intervals and respond to requests of the least popular content. These may shift their loads as peak user demand shifts over the course of the day across different geographic areas.

When Large Firms Operate Their Own Internet Infrastructure

Many large firms in application markets and platforms, such as Microsoft, Apple, Alphabet, Amazon, and Facebook, operate their own internet infrastructure rather than use third-party market suppliers. For example, all of them operate their own data centers and content delivery networks. As another example, Alphabet/Google connects its own data centers to each other with its own backbone lines and thus bypasses backbone lines it could lease from network operators. Large firms that integrate into complementary functions presumably do so because they can operate processes at a lower cost than third-party providers offer. It also may help achieve higher performance once the processes are tailored to specific needs. In the case of Google, for example, the lines help balance loads across its many data centers and CDNs over the course of the day. Large firms also may find scope economies across multiple related services, thereby spreading the efficiencies, or enabling them to offer services as bundled offerings that appeal to users (Bates et al. 2018). For example, as part of a suite of security offerings to protect content, Cloudflare offers CDN and name-server services inside one package of many services.

When large firms bring internet infrastructure in-house, the effects for the network as a whole can be positive, neutral, or even negative. For example, several of the largest firms that operate large data centers—like Microsoft, Google, and Amazon—began offering cloud services some years ago. Users became accustomed to the resulting efficiencies, and demand for these services is growing rapidly. The network economy, thus, benefited from the entry of these firms into the supply of cloud services.

However, the experience of Google Fiber illustrates another type of situation. Google started a new division to offer high-speed fiber to households and entered several cities with contracts for television, telephone, and internet service. While commercially successful in several cities, as of this writing, this division has paused its investments while seeking to overcome some challenges.[2] So far, therefore, the visible gains have been modest and localized to the few places where entry has been built or, at best, demonstrative of what might be possible elsewhere. Even if Google Fiber does cover all its intended cities, it will cover no more than 10 percent of the US population.

When large firms integrate into internet infrastructure, some outcomes of potentially greater concern arise; providers of complementary services must negotiate with large dominant firms, and thus potentially face contract terms they would not have encountered in a competitive setting with more options (Rogerson 2018). Also, there is a long-standing concern that increasing use of proprietary processes inside the largest firms can diminish the likelihood of generative innovations that

---

[2] As of this writing, Google Fiber offers service in Salt Lake City, NV; Provo, UT; Kansas City, MO; Austin, TX; Nashville, TN; Charlotte, NC; Atlanta, GA; Raleigh-Durham, NC; Orange County, CA; Huntsville, AL; and San Antonio, TX. It entered and exited Louisville, KY. Google Fiber has plans and permits to enter at least a dozen more cities, but no announced timeline.

would have arisen with wider use of open protocols (Zittrain 2008). These concerns play a significant role in antitrust or regulatory analysis. Therefore, an important open question for debate concerns the degree of market power and range of circumstances over which these concerns apply.

The rise of private data centers and the cloud, once again, raises important economic questions about competitive behavior and private incentives from improvements in networking infrastructure. What are the distribution of gains between users and producers from improvements in one part of a network, such as the cloud? In the presence of the gains shared by users, are competitive incentives sufficient? Do they favor some users over others? What are the long-term competitive prospects for new entry by entrepreneurial firms who use third-party services? These are important open research questions.

Protocols and Governance

Protocols are a set of rules and regulations that determine how data makes it through the network. A networking protocol defines conventions for processes, which include definitions for both the format of data packets, as well as for recovery in the event of transmission errors. For example, the TCP/IP (Transmission Control Protocol/Internet Protocol) is a family of protocols that sets a format for packets of data in the internet, defines addresses when networks connect to each other, defines how to assemble packets of data that arrive through the internet by different routes, and includes error correction processes. The BGP (Border Gateway Protocol) is the most commonly used protocol for routing traffic on the internet, although it is just one among many that governs how network switches and servers send packets of data through the network.

Engineers say equipment is compatible with other equipment only if both sets have adopted the same protocol. Each protocol lives with many complementary protocols in a protocol "stack"—a family of related protocols assembled together. The protocol stack acts as a reference model for designers, who largely aspire to make compatible equipment (for additional descriptions, useful starting points are Clark 2018; Knieps and Bauer 2016; Greenstein 2015).

The protocol stack for the internet (mostly) sends data packets along the least-congested route, a feature that delivers data quickly even when there are many potential routes for data and bottlenecked capacity along points of the network. This feature has become increasingly important because many modern internet applications, such as gaming and streaming, depend on fast delivery of data.

Infrastructure firms and carriers largely comply with protocol stacks; after all, they can offer profitable services while doing so. This outcome should not be taken for granted. It represents a notable departure from a prior era of practices, where many different firms offered proprietary protocols and networks and these did not interoperate. Since the birth of the commercial internet in the middle of the 1990s, however, compatibility with the internet protocol stack has been self-reinforcing. Compliance with protocols by all other suppliers and users further motivated widespread adoption and persistent use of these protocols by any participant, and it motivated development of many additional innovative services built on top of this equipment. The incentives for continuity are apparently strong in the modern internet, in spite of variance in the cause and size of the network effects across participants.

It is possible for situations to arise in which a break with the existing protocols makes sense to a decision-maker. Remarkably, none of those pressures has been sufficient in recent decades to generate stark breaks with internet protocols, though there are examples of partial movement in that direction (Simcoe and Watson 2019 provide a useful framework). For example, operators of the "dark web" prefer not make their content searchable, because they (allegedly) support illegal activities, such as the exchange of pirated material. Network effects also may not operate at the international level as different governments adopt mutually incompatible practices for their domestic networks, in some cases to censor content, but also to 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, packet-inspection processes in routers, or back-door design within operating systems to permit surveillance. These actions and policies frame open questions about the risks of losing seamless interoperability, or "splintering" the internet. These topics deserve attention from economic researchers.

These observations also motivate questions about the governance for improving protocols. For the most part, nonprofit organizations design and upgrade the protocol stack used for internet infrastructure. For example, the Internet Society provides the home for the Internet Engineering Task Force (IETF), which governs the protocols behind TCP/IP and BGP, and many others. The Internet Corporation for Assigned Names and Numbers (ICANN) governs assignment of domain names and updates the routing tables used by every switch and router on the internet. A routing table contains information about the topology of a network, and provides guidance about where data packets should go. In modern systems, the tables learn about congestion and send data on routes to avoid the congestion (for discussion, Clark 2018). The Institute of Electronic and Electrical Engineers (IEEE) convenes committee 802.11, which supports the standard underlying Wi-Fi, as well as other technical standards. These organizations convene groups that design, maintain, and upgrade the protocols, and they subsequently charge little for their use. Most also put few legal restraints on how the private sector operates the equipment that uses those protocols.

Many voices influence and determine the actions of the organizations who govern protocols. Given the private stakes, it is no surprise that debates about policies for intellectual property receive considerable attention today, as do debates for criteria about what administrative process should be used to create a protocol. For an example of such a debate, consider the problem of exhaustion of available

IP addresses, which necessitated a redesign of IP addresses to enable growth into the future. Version 6, abbreviated as IPv6, emerged from a debate at the Internet Engineering Task Force. It has been slow to diffuse since it became available. Many blame the new design, which is cumbersome to adopt.

As another example, consider the vociferous debate surrounding the expansion of top-level domain names by ICANN. The internet was designed with 248 country codes, but six domains inside the United States, where no country code was required—com, org, net, edu, gov, and mil—became widely used, especially com. In response to complaints about the limitations arising from the concentration of names under com, ICANN expanded the number of domains to over 1,000, including icu, top, xyz, site, vip, and online. For histories of these and related organizations, and an analysis of their origins, interested readers might begin with Mueller (2004), Simcoe (2012), Russell (2014), Greenstein (2015), and Clark (2018).

The choice of protocols and changes to protocols resembles a public good problem because virtually all users have the same experience, and they can neither opt out nor be excluded from changes. Considerable effort goes into the design of protocols, but not all of them receive equal use. Development of economic theory for when it is worthwhile to change protocols gradually or dramatically, or whether to abandon them at all, is essential for understanding the continuance of the commercial internet.

Pricing and Incentives

For most internet users, the up-front price they face involves a fee from their internet service provider. The vast majority of business users in urban and suburban areas contract for broadband internet access (for the diffusion of broadband, see Ryan and Lewis 2017, Pew Research Center 2019). From 2012 to 2017, payments for access to wireline forms reached $88.7 billion, growing more than 30 percent. Wireline access also became faster, as much as doubling in speed between 2011 and 2018. Payments for access fees to wireless service also reached over $90 billion, an increase of 57 percent (according to Census data from the Statistics of US Business). The revenue increase during this period did not largely arise from a rise in the number of households because most US households already had internet service in 2012.[3]

The market for supply of broadband services has a moderate degree of competition. This supply structure emerged after the replacement of dial-up with broadband as the primary method of internet access (Greenstein and McDevitt 2011). In 1995, virtually all internet access occurred over dial-up; whereas today, approximately 80 percent of US households have broadband internet access in

---

[3] For example, the Netflix ISP Speed Index comparisons of measured speeds over 2012-2018 yields a doubling of realized speeds for most networks (https://ispspeedindex.netflix.com/country/us/). From 2011 to 2018, only 3 to 5 percent of US households first began using broadband internet, depending on the survey (Pew Research Center 2019).

their homes. Downtown locations in high-density settings experienced greater entry, aimed at business customers and/or multi-occupation residences (Chen and Savage 2011; Connolly and Preiger 2013). Most households in urban and suburban settings have access to at least one or two providers of wireline access, and multiple wireless providers (Wallsten and Mallahan 2013). The typical providers for households are the local telephone company, who typically offers DSL service, and the local cable television provider, who offers data services using modems compatible with DOCSIS, the Data Over Cable Service Interface Specification. In some areas, a third-party "over-builder" may offer fiber to the home, and one local telephone company, Verizon, offers fiber to homes in some of the territories in which it operates. Business in dense urban locations may have access to even more providers.

Some of the many components that play a role in limiting entry of internet access providers include financial reasons, such as high capital costs; regulatory factors, such as limited rights of way, rules raising the costs of over-builders, and laws preventing entry from municipal providers (for example, Seamans 2012); and behavioral dynamics, such as the unwillingness of incumbent firms to enter each other's established territories. In addition, technical forces make some forms of broadband access, such as DSL or 5G wireless service, less effective outside of dense locations (for example, Destafano, Kneller, and Timmis 2018), or make cable service cost-prohibitive. Satellite services remain viable in most terrain, providing a baseline level of service for less-dense areas (Boik 2017). Relatedly, a robust market for supplying cellular towers to support carriers' antennae enables service from two to four providers in all but the least dense locations.

Meanwhile, measured price levels for access have changed little since broadband became the dominant delivery mode for households. The Consumer Price Index (available at US Bureau of Labor Statistics 2020) provides a measure of broadband prices in the price series for "Internet services and electronic information providers" (US city average, all urban consumers), which rises from 73.4 to 77.1 from 2007 to 2019, an increase of 5 percent. The closest comparable index for wireless services (which covers data and also includes the price of telephone calls) shows the Consumer Price Index for "Wireless telephone services" (again, US city average, all urban consumers) dropped from 64.5 to 46.4 over the same period, a decline of 22 percent. In light of the enormous changes in those years—for example, the rise of Web2.0 businesses, the growth of social media, and the explosion of short and long form video and streaming—it is likely that measured prices miss important aspects of the typical user experience.

What is missing? For one, neither index adjusts the price of internet access for the quality of that service. In addition, neither accounts for changes in the quality of ad-supported "free" content (for an approach to the latter, see Byrne and Corrado 2019). Lastly, large growth in access revenues with small growth in the number of subscribers indicates many households increased their expenditure on internet access by moving to higher tiers of service at higher prices. Standard methods for price measurement do not count such migration across tiers necessarily as a change in prices.

When do wireline and wireless services substitute for each other, and when do they complement each other? No general answer exists. Any answer changes over time as access capabilities improve and as modal applications change, and it varies by location of the user. In some applications today, wireline and wireless delivery, such as electronic mail and passive browsing, can substitute when users can tolerate delays. These modes of delivery do not substitute in other applications, such as data-intensive streaming and gaming, where delays interfere with user experience. Sometimes they complement each other, such as when entertainment firms encourage tweeting during an online gaming event or streaming of content. These are difficult questions for a substantial number of households that get their internet through only a wireless smart phone or satellite. As access to frontier infrastructure improves, the extent of substitution and complementarity between wireless and wireline services is an important open topic of research.

Contracts between users and access firms also changed over time. In the earliest years, most access involved a monthly charge and no limitations on use. Greenstein (2015) discusses the disappearance of price discrimination based on the amount of time online. Today price discrimination based on usage of data, combined with data caps, is common in both wireline and wireless contracts. Moreover, wireline and wireless data contracts do not take the same form. Burnham et al. (2013) provides early census of the use of tiered pricing and caps based on the usage of data in wireline. Recent studies show that some users are sensitive to the charges affiliated with reaching a data cap, but they also endogenously select into capacity consistent with their use; for example, those who practice heavy streaming choose plans that allow this without large cost increases (for example, Nevo, Turner, and Williams 2016; McManus et al. 2018).

The incentives for improved internet capabilities and access receive considerable attention from policy analysts. On the one hand, there is the general belief that improvements in the speeds of wireline and wireless access benefit more participants than just the firm providing this access. While access providers potentially gain more revenue, users gain better service, and application providers face the option of a new frontier for their data-intensive services. Once again, the gains from improvement are widespread, while the private costs and commercial risks are concentrated in the one investor—in this example, access providers. As mentioned earlier, an important open question concerns the gaps between private and social incentives to upgrade wireline broadband. Nevo, Turner, and Williams (2016) suggest the gap is substantial, consistent with the presence of insufficient private incentives to upgrade quality at a rate in line with society's broader interest.

Estimating these incentives remains an open research area, especially in upgrades to wireless technology. For example, cellular telephony migrated to new generations of technologies, from 3G to 4G. 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 summary, users have increased the use of data substantially on wireless modes as it has deployed, and suppliers have invested in order to support those increased volumes. Were incentives for this evolution too high or too low, and how would answers to this question inform expectations about the ongoing upgrade to 5G?

Another quality-related concern touches on competitive issues. If wireline broadband firms carry video-on-demand, then questions arise over conflict of interest in carrying other forms of internet traffic, which in turn would effect investment, interconnection, and pricing (Rogerson 2018). One other change also may have shaped incentives in the recent experience, and may do so in the future. At the outset of the commercial internet, internet service providers did not charge for accepting data delivered to them to be sent to their direct customers, but today some do. This provides additional revenue for internet service providers, and it has been met with resistance from other providers who interconnect with ISPs because it raises the costs of providing data services over long distances and content delivery networks.

Less data is available concerning the prices and fees that happen behind the scenes in interconnection. Interconnection networks often reach agreement without disclosing terms. Evidence from Zhuo et al. (2019) shows interconnection agreements growing everywhere, with some variance across different geographies due to economic development. Negotiation plays an important role in shaping fees for private peering and content delivery networks. Any time that price negotiations take place in the shadow of alternative options for accomplishing functionally equivalent outcomes, then those options discipline attempts to raise prices or exploit negotiating advantage in other ways. Conversely, infrastructure firms have a negotiating advantage when they provide services for which there are no substitutes, and/ or when they can bottleneck the aspirations of other network participants.

In this setting, are prices inside the internet infrastructure more likely to be determined by a range of competitive options that tend to drive down prices paid by ultimate users? Or are these prices more likely to be determined in many situations of limited competition and bottlenecks? There is limited evidence on these questions. An optimistic view focuses on how much all participants know about the conduct of negotiations (Norton 2014). Often negotiations resolve themselves without incident: in fact, there has not been a prominent example of breakdown in negotiations since late 2013, when Netflix and the four largest providers of internet access in the United States could not reach a negotiated settlement, which led to widespread congestion issues affected streaming speeds and reliability at tens of millions of households (Greenstein and Norris 2015). The underlying issue in this case was that the original model for the commercial internet involved no charges for delivery of data to internet service providers, and attempts to impose such charges played a role in the negotiation breakdown between Netflix and four large ISPs.

Perhaps the most salient evidence for being optimistic that these prices are being shaped and determined by a range of (reasonably) competitive options is the long-term record of a symbiotic relationship between advances in infrastructure, growth in access revenue, and advances in revenue for electronic commerce. Internet infrastructure has contributed to reducing several frictions related to conducting commerce; indeed, Goldfarb and Tucker (2019) argue that digital technology has improved economic activity largely through reduction in these frictions. The earlier illustration of a transaction focused on a user's request for data from Wikipedia, a nonprofit organization, but if the user had requested data from profit-seeking firms, for example, it would have triggered additional commercial actions. If there had been advertising, then advertisers would have paid for ad exchanges and geolocation of the IP address, so the user receives a geographically-appropriate advertisement. Related processes may have personalized the advertising further. All of these steps would take place virtually instantly, and largely unseen to the user. Had the user bought or sold a product, many additional processes would have supported fulfillment of the order and would have increased the flow of funds to infrastructure—data centers, switches, and transmission lines—to support the transactions.

A pessimistic view of the internal pricing of the internet points out that the features of any known incident lack transparency, especially in the first decade after the millennium (Greenstein 2010). In the 2013 Netflix incident, for example, households did not know who unreasonably held up whom—the local access provider or Netflix—as performance declined for many applications other than those involved in negotiations. Because the terms for settlement did not become public, competing interpretations of the event remain unresolved, and so too do questions about whether government intervention could have alleviated the decline in internet performance experienced by users.

A pessimist also would observe how the dependence of network participants on many firms creates difficulties in assigning responsibility for inadequacies in the delivery of service. An outage can have widespread consequences. For example, when services like Slack, Quora, and Medium all became unavailable February 28, 2017, users had no way to know that Amazon cloud storage had gone down due to a single maintenance person's faulty actions at one AWS facility in Northern Virginia, which caused a number of Amazon web servers to go offline. As another example, when services such as CNBC, Netflix, and Twitter went down on October 21, 2016, users had no way to know that it resulted from a distributed denial-of-service attack on Dyn, who provided the name-server services for these firms. However, events like these often provide the fodder for debate about legal or regulatory frameworks for denoting who assumes responsibility for compensation from economic loss, and relatedly, whether these frameworks provide sufficient incentives to suppliers for investment in risk mitigation.

Geographic Availability

An uneven geographic supply of internet infrastructure is not the only reason why some areas have high rates of non-adoption of internet access, but it plays a major role. Today approximately 10 percent of the US population does not use the internet (Anderson et al. 2019). Some of that non-adoption is linked to demographic features of users, such as older age, low income, and less education. But an important factor is the location of a household, namely, in a rural or low-density location. While 97 percent of the land in the United States is rural, according to the Census Bureau, 19 percent of the population lives in rural locations—that is, areas with sparse residential housing.

Cutting-edge internet infrastructure tends not to be available in low-density regions. In some of these locations, even internet infrastructure with older technology may not be available (for additional discussion, see Forman et al. 2018). For example, according to a Pew survey conducting in February 2019, 63 percent of US rural residents say they have broadband access, compared with 80 percent for suburban residents (Perrin 2019). Moreover, 20 percent of rural households are wireless-only; for comparison, 25 percent of low-income households are wireless-only (Anderson 2019). "Unavailable broadband service" or "low quality" is cited by 22 percent of non-broadband home users as reasons for relying solely on wireless smartphones for internet service, among those who do. However, the price of wireline broadband access is the most frequently cited reason to go to only wireless, with more than half the users citing high prices as the primary reason.

An uneven geographic supply of internet infrastructure arises for many reasons. The costs of supplying internet service to a given geographic areas may reflect economies of scale: that is, when installing and operating cell towers, data centers, and content delivery networks, these structures endogenously locate near a greater number of densely located users, because it facilitate faster returns on investment. Laying lines between locations involves high fixed costs and low marginal costs, so low density may not have sufficient demand to incentive such investments. A demand for higher quality can also drive unequal dispersion because suppliers prefer to build out initially in more affluent and urban locations where a greater number of buyers are more willing to pay for the expensive frontier quality. Marshallian agglomeration can reinforce these differences, with richer, dense urban locations attracting skilled labor and receiving infrastructure closer to the technological frontier.

Geographic variance in supply potentially creates different experiences across households and businesses in different locations. Econometricians often look for this type of variance. However, many variations arise only at a fine level of geography, such as a neighborhood, and attempts to measure availability at this fine-grained level have encountered numerous challenges. For example, an attempt to create a National Broadband Map, which began in 2011, went through several revisions; it was regarded as accurate in some but not all locations, and was discontinued in December 2018. As of this writing, the Federal Communications Commission is developing a new mapping program.

Public policy in such situations also faces some tradeoffs affiliated with economic factors that defy easy measurement. Many users prefer a local supply of internet infrastructure when it is available. However, it may be cheaper to use remote data centers, cloud storage and/or satellites, and users may be willing to substitute into distant suppliers under a range of circumstances. These consumer preferences are challenging to learn, but they will shape any calculation of how much society is willing to spend to provide basic internet access to areas where it would not otherwise be supported by the market, and how much society is willing to spend to provide a quality of internet service to these areas above the basic level.

Many different programs seek to address these concerns. For example, the 1996 Telecom Act established the e-rate program, which taxed telephone calls to finance subsidies for rural broadband. Today it raises more than $4 billion 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 low-density areas. Programs to address demand also exist, but are less common. As a condition for approval of a merger, for example, Comcast agreed to offer lower prices to qualifying low-income households, and evidence suggests it had an effect on hundreds of thousands of households (Rosston and Wallsten 2019). Given the range of programs, it is no surprise that there are many debates over the effectiveness of subsidies of different sizes and designs.

Final Policy Questions

Internet infrastructure has been improving over the decades in ways that have enabled an extraordinary gain in internet services, which many users were willing to pay for. In turn, this reinforced incentives to do something less visible to most users—namely, to improve digital infrastructure. The virtuous cycle has gone hand-in-hand with growth in access revenues, growth in advertising revenues for free services (such as search and news), and growth in electronic commerce that takes advantage of online shopping.

A big, open question is whether this growth and improvement will continue into the future. Some factors will slow growth, such as saturated adoption of broadband in households and businesses. Other factors may accelerate it, such as the restructuring of online services to take advantage of the cloud and 5G wireless infrastructure. The history of the last few decades suggests that internet architecture contains a remarkable capacity to adapt to changes but, in fact, every change raises novel challenges. Every change alters revenues and costs for different suppliers, and some of these developments generates disputes, such as those that accompanied the early diffusion of streaming. So it is an open question whether certain kinds of future changes will place undue stress on the governance of internet infrastructure and the resilience of its designs.

So far, no consensus has emerged regarding an internet regulatory framework. Instead, policy has evolved alongside commercial internet growth, including formal policy in statute or regulatory orders (like in conditions for mergers), less formal policy found in public statements about general principles (like speeches from the chair of the Federal Communications Commission), along with the willingness of government actors to intervene. For a history of these policies, see Nuechterlein and Weiser (2005), Greenstein (2010), Greenstein, Peitz, and Valleti (2016), Knieps and Bauer (2016), and Cybertelecom.org.

No futurist foresees a lack of opportunity to restructure wireless and cloud services, nor does anyone foresee ubiquitous competitive broadband arising in all locations in the next few years. These and other internet changes are sure to generate open questions and policy debate for the foreseeable future.

■ *The author would like to thank Tim Bresnahan, David Clark, kc Claffy, Chris Forman, Avi Goldfarb, Gordon Hansen, John Laprise, Enrico Moretti, Abhishek Nagaraj, Frank Nagle, Tommy Pan Fang, Bill Rogerson, Alicia Shems, Heidi Williams, and Pai-Ling Yin for extremely helpful conversations and comments. Timothy Taylor deserves special thanks for masterful editing.*

References

Anderson, Monica. 2019. "Mobile Technology and Home Broadband 2019." *Pew Research Center,* June 13. https://www.pewresearch.org/internet/2019/06/13/mobile-technology-and-home-broadband-2019/.

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 22. https://www.pewresearch.org/fact-tank/2019/04/22/some-americans-dont-use-the-internet-who-are-they/.

Bates, Samantha, John Bowers, Shane Greenstein, Jordi Weinstock, and Jonathan Zittrain. 2018. "In Support of Internet Entropy: Mitigating an Increasingly Dangerous Lack of Redundancy in DNS Resolution by Major Websites and Services." NBER Working Paper 24317.

Boik, Andre. 2017. "The Economics of Universal Service: An Analysis of Entry Subsidies for High Speed Broadband." *Information Economics and Policy* 40: 13–20

Burnham, Brad, Shane Greenstein, Neil Hunt, Kevin McElearney, Marc Morial, Dennis Roberson, and Charles Slocum. 2013. *Issues in Data Caps and Usage Based Pricing.* Washington, DC: Federal Communications Commission.

Byrne, David, and Carol Corrado. 2019. "Accounting for Innovation in Consumer Digital Services: IT Still Matters." NBER Working Paper 26010.

Byrne, David, Carol Carrado, and Dan Sichel. 2018. "The Rise of Cloud Computing: Minding Your P's, Q's, and K's." NBER Working Paper 25188.

Chen, Yongmin, and Scott J. 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.

Connolly, Michelle, and James E. Prieger. 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." Economic Statistics Centre of Excellence Discussion Paper 2018-2019.

DeStafano, Timothy, Richard Kneller, and Jonathan Timmis. 2018. "Broadband Infrastructure, ICT Use, and Firm Performance: Evidence for UK Firms." *Journal of Economic Behavior and Organization* 155: 110–39.

Ewens, Michael, Ramana Nanda, and Matthew Rhodes-Kropf. 2019. "Cost of Experimentation and the Evolution of Venture Capital." Harvard Business School Working Paper 15-070.

Federal Communications Commission. 2013. *Measuring Broadband America—2014.* Washington, DC: Federal Communications Commission.

Federal Communications Commission. 2018. *Measuring Fixed Broadband-Eighth Report.* https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-fixed-broadband-eighth-report. Accessed March, 2020.

Feldstein, Martin. 2017. "Underestimating the Real Growth of GDP, Personal Income, and Productivity." *Journal of Economic Perspectives* 31 (2): 145–64.

Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57 (1): 3–43.

Greenstein, Shane. 2010. "Glimmers and Signs of Innovative Health in the Commercial Internet." *Journal of Telecommunication and High Technology Law* 8 (1): 25–78.

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

Greenstein, Shane, Chris Forman, and Avi Goldfarb. 2018. "How Geography Shapes—and is Shaped by—the Internet." In *The New Oxford Handbook of Economic Geography*, edited by Gordon L. Clark, Maryann P. Feldman, Meric S. Gertler, and Dariusz Wojcik, 269–85. Oxford, UK: Oxford University Press.

Greenstein, Shane, Martin Peitz, and Tomasso Valleti. 2016. "Net Neutrality: A Fastlane to Understanding the Tradeoffs." *Journal of Economic Perspectives* 30 (2): 127–50.

Greenstein, Shane, and Ryan C. McDevitt. 2011. "The Broadband Bonus: Estimating Broadband Internet's Economic Value." *Telecommunications Policy* 35 (7): 617–32.

Greenstein, Shane, and Michael Norris. 2015. "Streaming Over Broadband: Why Doesn't My Netflix Work?" Harvard Business School Case 616-007.

Groshen, Erica L., Brian C. Moyer, Ana M. Aizcorbe, Ralph Bradley, and David M. Friedman. 2017. "How Government Statistics Adjust for Potential Biases from Quality Change and New Goods in an Age of Digital Technologies: A View from the Trenches." *Journal of Economic Perspectives* 31 (2): 187–210.

Huston, Geoff. 2017. *The Rise and Rise of Content Distribution Networks.* San Diego, CA: Center for Applied Internet Data Analysis.

Internet and Television Association. 2018. "Wi-Fi: How Broadband Households Experience the Internet." *Internet and Television Association*, April 6. https://www.ncta.com/whats-new/wi-fi-how-broadband-households-experience-the-Internet.

Knieps, Günter, and Johannes M. Bauer. 2016. "The Industrial Organization of the Internet." In *Handbook on the Economics of the Internet*, edited by Johannes M. Bauer and Michael Latzer, 23–54. Cheltenham, UK: Edward Elgar Publishing.

McManus, Brian, Aviv Nevo, Zachary Nolan, and Jonathan W. Williams. 2018. "Steering Incentives and Bundling Practices in the Telecommunications Industry." Networks, Electronic Commerce, and Telecommunications Institute Working Paper 18-12.

Mueller, Milton. 2004. *Ruling the Root: Internet Governance and the Taming of Cyberspace.* Cambridge, MA: MIT Press.

Nakamura, Leonard, Jon Samuels, and Rachel Soloveichik. 2016. "Valuing 'Free Media' in GDP: An Experimental Approach." Federal Reserve Board of Philadelphia Working Paper 16-24.

Nevo, Aviv, John L. Turner, and Jonathan W. 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 Phillip Weiser. 2005. *Digital Crossroads: Telecommunications Law and Economics in the Internet Age.* Cambridge, MA MIT Press.

Perrin, Andrew. 2019. "Digital Gap Between Rural and Non-rural Persists." *Pew Research Center*, May 31. https://www.pewresearch.org/fact-tank/2019/05/31/digital-gap-between-rural-and-nonrural-america-persists/.

Pew Research Center. 2019. *Internet/Broadband Fact Sheet.* Washington, DC: Pew Research Center. https://www.pewresearch.org/internet/fact-sheet/internet-broadband/.

Pew Research Center. 2019. *Mobile Fact Sheet.* Washington, DC: Pew Research Center. https://www.pewresearch.org/internet/factsheet/mobile/.

Rogerson, William P. 2018. "Economic Theories of Harm Raised by the Proposed Comcast/TWC

Transaction." Chapter 18 in *The Antitrust Revolution,* 7th ed., edited by John E. Kwoka, Jr. and Lawrence J. White. Oxford, UK: Oxford University Press.

Rosston, Gregory L., and Scott J. Wallsten. 2019. "Increasing Low-Income Broadband Adoption through Private Incentives." Unpublished. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3431346.

Russell, Andrew. 2014. *Open Standards and the Digital Age: History, Ideology, and Networks.* Cambridge, UK: Cambridge University Press.

Ryan, Camille, and Jamie M. Lewis. 2017. *Computer and Internet Use in the United States: 2015.* Washington, DC: US Census Bureau.

Seamans, Robert C. 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.

Simcoe, Timothy, and Jeremy Watson. 2019. "Forking, Fragmentation, and Splintering." *Strategy Science* 4 (4): 283–97.

Syverson, Chad. 2017. "Challenges to Mismeasurement Explanations for the US Productivity Slowdown." *Journal of Economic Perspectives* 31 (2): 165–86.

US Bureau of Labor Statistics. 2020. *Consumer Price Index for All Urban Consumers. US City Average. Internet Services and Electronic Information Providers.* https//data.bls.gov/PDQWeb/cu.

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, 131–56. Cheltenham, UK: Edward Elgar Publishing.

Wang, Jin, and Kristina Steffenson McElheran. 2017. "Economies before Scale: Survival and Performance of Young Plants in the Age of Cloud Computing." Rotman School of Management Working Paper 3112901.

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.

Zittrain, Jonathan. 2008. *The Future of the Internet and How to Stop it.* New Haven, CT: Yale University Press.