

Oxford Handbooks Online

Internet Infrastructure

Shane Greenstein

The Oxford Handbook of the Digital Economy

Edited by Martin Peitz and Joel Waldfogel

Print Publication Date: Aug 2012 Subject: Economics and Finance, Economic Development

Online Publication Date: Nov 2012 DOI: 10.1093/oxfordhb/9780195397840.013.0001

Abstract and Keywords

This article presents an overview on Internet infrastructure, highlighting Internet access and broadband development. Internet infrastructure is a collective term for all the equipment, personnel, and organizations that support the operation of the Internet. The Internet largely used existing capital at many installations with comparatively minor software retrofits, repurposing capital for any new function as long as it remained consistent with end-to-end. As Internet applications became more popular and more widely adopted, users pushed against the bandwidth limits of dial-up Internet access. Broadband gave users a better experience than dial-up access. The most notable feature of the governance of Internet infrastructure are the differences with the governance found in any other communications market. Many key events in Internet infrastructure took place within the United States, but this seems to be less likely as the commercial Internet grows large and more widespread.

Keywords: Internet infrastructure, Internet access, broadband, United States, governance

1. Introduction

Internet infrastructure is a collective term for all the equipment, personnel, and organizations that support the operation of the Internet. The commercial Internet arose after many firms and users voluntarily adopted a set of practices for enabling internetworking, namely, transferring data between local area networks and computer clients. The commercial Internet began to provide many revenue-generating services in the mid-1990s. As of this writing, this network supports a wide array of economic services to more than a billion users, and it continues to grow worldwide.

Generally speaking, four types of rather different uses share the same Internet infrastructure: browsing and e-mail, which tend to employ low bandwidth and which can tolerate delay; video downloading, which can employ high bandwidth and can tolerate some delay; voice-over Internet protocol (IP) and video-talk, which tend to employ high bandwidth and whose quality declines with delay; and peer-to-peer applications, which tend to use high bandwidth for sustained periods and can tolerate delay, but, in some applications (such as Bit-Torrent), can impose delay on others.¹

The precise economic characteristic of Internet infrastructure defies simple description. Like other private assets, sometimes Internet infrastructure is an input in the production of a pecuniary good, and regular investment extends the functionality or delays obsolescence. Like a public good, sometimes more than one user employs Internet infrastructure without displacing another or shaping the (p. 4) quality of another user's experience. Even when visible, the economic contribution of Internet infrastructure is often measured indirectly at best. Yet in contrast to many private assets or public goods, sometimes many decision makers—instead of one—govern the creation and deployment of Internet infrastructure. Moreover, although market-oriented prices influence the investment decisions of some participants, something other than market prices can shape the extent of investment and the use of such investment.

This array of applications, the mix of economic characteristics, and the economic stakes of the outcomes have attracted considerable attention from economic analysts. This review summarizes some of the key economic insights of the sprawling and vast literature on Internet infrastructure.

The chapter first provides a summary of the emergence of the commercial Internet, providing a review of the common explanations for why its structure and pricing took a specific form. The chapter then describes how the deployment of broadband access and wireless access altered many aspects of Internet infrastructure. It also reviews the importance of platforms, another new development that has changed the conduct of firms providing infrastructure for the commercial Internet. The chapter finishes with a review of the private governance of Internet infrastructure and the role of key government policies.

2. The Emergence of the Commercial Internet

Which came first, the infrastructure or the commercial services? In the popular imagination the Internet emerged overnight in the mid-1990s, with the creation of the commercial browser and services employing that browser. In fact, the Internet developed

for more than two decades prior to the emergence of the browser, starting from the first government funding at the Defense Advanced Research Project Agency (DARPA) in the late 1960s and the National Science Foundation (NSF) in the 1980s.² Hence, there is no question that the infrastructure came first. The Internet infrastructure of the mid-1990s employed many components and inventions inherited from the government-funded efforts.

Both the commercial Internet and its government-sponsored predecessor are packet switching networks.³ While the Internet is certainly not the first packet switching network to be deployed, the commercial Internet has obtained a size that makes it the largest ever built.

In a packet switching network a computer at one end packages information into a series of discrete messages. Each message is of a finite size. As part of initial processing by the packet switching system, larger messages were divided into smaller packets. All of the participating computers used the same conventions for messages and packets. Messages successfully traveled between computers when the (p. 5) sending and receiving computer standardized on the same procedures for breaking up and reassembling packets.

Sending data between two such locations rely on something called “protocols,” standardized software commands that organize the procedures for moving data between routers, computers, and the various physical layers of the network. Protocols also define rules for how data are formatted as they travel over the network. In 1982 one design became the standard for the network organized by DARPA, a protocol known as TCP/IP, which stands for transmission control protocol/internet protocol. Vint Cerf and Robert Kahn wrote the first version, replacing an outmoded protocol that previously had run the DARPA network. Over time the technical community found incremental ways to change the protocol to accommodate large-scale deployment of infrastructure and applications.⁴

TCP/IP had the same role in the government-sponsored and commercial Internet. It defined the “headers” or labels at the beginning and end of a packet.⁵ Those headers informed a computer processor how to reassemble the packets, reproducing what had been sent. By the mid-1980s all Unix-based computing systems built that interface into their operating systems, and in the 1990s virtually every computing system could accommodate it.⁶

An example of TCP/IP in action may help to illustrate its role. Consider sending messages from one computer, say, at Columbia University in New York, to another, say, at Stanford University in California. In the 1980s such an action involved the following steps. A user of a workstation, typically using a computer costing tens of thousands of dollars, would

compose the message and then issue a command to send it. This command would cause the data to be broken into TCP/IP packets going to a central mail machine at Columbia and then a destination mail machine at Stanford. In between the central mail machine and destination mail machine the file would be broken down into pieces and completely reassembled at the final destination. During that transmission from one machine to another, the file typically traveled through one or more routers, connecting various networks and telephone lines. At the outbound site it was typically convenient to send through a central mailserver, though that did not need to be true at the inbound site. At the Stanford server, the local computer would translate the data into a message readable on another person's workstation. If something went wrong—say, a typographical error in the address—then another set of processes would be initiated, which would send another message back to the original author that something had failed along the way.⁷

These technical characteristics define one of the striking economic traits of Internet infrastructure: it requires considerable coordination. First, designers of hardware and software must design equipment to work with other equipment, and upgrades must meet a similar requirement. Second, efficient daily operations of the Internet require a wide array of equipment to seamlessly operate with one another. Indeed, the widespread adoption of TCP/IP partly explains how so many participants coordinate their activity, as the protocol acts as a focal point around which many participants organize their activities at very low transaction costs. (p. 6) Another key factor, especially during this early period, helped coordinate activity, a principle called *end-to-end*.⁸ This principle emerged in the early 1980s to describe a network where the switches and routers retained general functionality, moving data between computers, but did not perform any processing. This was summarized in the phrase, “The intelligence resided at the edges of the network.” Any application could work with any other at any edge, so long as the routers moved the data between locations.⁹

End-to-end had substantial economic consequences for the deployment of the commercial Internet. At the time of its invention, it was a radical engineering concept. It differed from the principles governing the telephone network, where the switching equipment performed the essential processing activities. In traditional telephony, the end of the network (the telephone handset) contained little functionality.¹⁰ For a number of reasons, this departure from precedent was supported by many outsiders to the established industry, especially researchers in computing. Hence, the Internet encountered a mix of benign indifference and active resistance from many established firms in the communications market, especially in the telephone industry, and many did not invest in the technology, which later led many to be unprepared as suppliers to the commercial Internet in the mid-1990s.¹¹

End-to-end also had large economic consequences for users. First, the Internet largely employed existing capital at many installations with comparatively minor software retrofits, repurposing capital for any new function as long as it remained consistent with end-to-end. Repurposing of existing capital had pragmatic economic advantages, such as lowering adoption costs. It also permitted Internet applications and infrastructure to benefit from the same technical changes shaping other parts of computing.¹²

Accommodating heterogeneous installations supported another—and, in the long run, especially important—economic benefit: It permitted users with distinct conditions to participate on the same communications network and use its applications, such as email.

3. The Structure of the Commercial Internet

As long as the Internet remained a government-managed enterprise, the NSF-sponsored Internet could carry only traffic from the research community, not for commercial purposes.¹³ In the late 1980s the NSF decided to privatize the operations of a piece of the Internet it managed (while the military did not privatize its piece). One key motivation for initiating this action was achievement of economies of scale, namely, the reduction in costs that administrators anticipated would emerge if researchers shared infrastructure with private users, and if competitive processes drove private firms to innovate.¹⁴ The transition reached resolution by the mid-1990s. Many of the events that followed determined the structure of (p. 7) supply of Internet infrastructure for the next decade and a half, particularly in the United States.¹⁵

The first group of backbone providers in the United States (i.e., MCI, Sprint, UUNET, BBN) had been the largest carriers of data in the NSF network. In 1995 and 1996, any regional Internet service provider (ISP) could exchange traffic with them. At that time the backbone of the US Internet resembled a mesh, with every large firm both interconnecting with every other and exchanging traffic with smaller firms.¹⁶ Some of these firms owned their own fiber (e.g., MCI) and some of them ran their backbones on fiber rented from others (e.g., UUNet).

After 1997 a different structure began to take shape. It was partly a mesh and partly hierarchical, using “tiers” to describe a hierarchy of suppliers.¹⁷ Tier 1 providers were national providers of backbone services and charged a fee to smaller firms to interconnect. The small firms were typically ISPs that ranged in size and scale from wholesale regional firms down to the local ISP handling a small number of dial-in customers. Tier 1 firms did most of what became known as “transit” data services, passing data from one ISP to another ISP, or passing from a content firm to a user. In

general, money flowed from customers to ISPs, who treated their interconnection fees with backbone firms as a cost of doing business.

Tier 1 firms adopted a practice known as “peering,” and it appeared to reinforce the hierarchical structure. Peering involved the removal of all monetary transfers at a point where two tier 1 providers exchanged traffic. Peering acknowledged the fruitlessness of exchanging money for bilateral data traffic flows of nearly equal magnitude at a peering point. Hence, it lowered transaction costs for the parties involved. However, because their location and features were endogenous, and large firms that denied peering to smaller firms would demand payment instead, many factors shaped negotiations. As a result, the practice became controversial. An open and still unresolved policy-relevant economic question is whether peering merely reflects a more efficient transaction for large-scale providers or reflects the market power of large suppliers, providing additional competitive advantages, which smaller firms or non-tier 1 firms cannot imitate. The economics of peering is quite challenging to characterize generally, so this question has received a range of analysis.¹⁸

Another new feature of Internet infrastructure also began to emerge at this time, third-party caching services. For example, Akamai, a caching company, would pay ISPs to locate its servers within key points of an ISP's network. Content providers and other hosting companies would then pay the caching companies to place copies of their content on such services in locations geographically close to users, aspiring to reduce delays for users. Users would be directed to the servers instead of the home site of the content provider and, thus, receive faster response to queries. A related type of service, a content delivery network (CDN), provided similar services for electronic retailers. In general, these became known as “overlays,” and since these were not part of the original design of the noncommercial Internet, numerous questions emerged about how overlays altered the incentives to preserve end-to-end principles.¹⁹

These changes also raised an open question about whether a mesh still determined most economic outcomes. For much Internet service in urban areas the (p. 8) answer appeared to be yes. ISPs had options from multiple backbone providers and multiple deliverers of transit IP services, and many ISPs multihomed to get faster services from a variety of backbone providers. ISPs also had multiple options among cache and CDN services. Evidence suggests that prices for long-distance data transmission in the United States continued to fall after the backbone privatized, reflecting excess capacity and lower installation costs.²⁰

For Internet service outside of urban areas the answer appeared to be no. ISPs did not have many options for “middle mile” transit services, and users did not have many

options for access services. The high costs of supply made it difficult to change these conditions.²¹

Perhaps no factor altered the structure of supply of Internet infrastructure after commercialization more than the World Wide Web: The World Wide Web emerged just as privatization began, and it is linked with a particularly important invention, specifically, the commercial browser. As the commercial Internet grew in the mid to late 1990s, traffic affiliated with the web overtook electronic mail, symptomatic of web applications as the most popular on the Internet. The growth in traffic had not been anticipated at the time NSF made plans for privatizing the backbone, and the subsequent growth in traffic fueled a global investment boom in Internet infrastructure.

Tim Berners-Lee built key parts of the World Wide Web.²² In 1991 he made three inventions available on shareware sites for free downloading: html (hypertext markup language) and the URL (universal resource locator), a hyper-text language and labeling system that made transfer of textual and nontextual files easier using http (hyper-text transfer protocol). By 1994 and after the plans for commercialization were set and implemented, a new browser design emerged from the University of Illinois. It was called Mosaic, and it was widely adopted in the academic Internet. Mosaic became the model that inspired the founding of Netscape, the first successful commercial browser, which also motivated Microsoft to develop Internet Explorer, which touched off the browser wars.²³

As of this writing most observers expect data traffic to continue to grow at rates in the neighborhood of 40 percent to 50 percent a year.²⁴ This is due to the emergence of another set of applications. Data supporting peer-to-peer and video applications have been growing rapidly after the millennium, each year using larger fractions of available capacity.²⁵

4. Pricing Internet Access

The value chain for Internet infrastructure begins with the pricing for Internet access. Households and business establishments pay for access to the Internet.²⁶ ISPs provide that access and use the revenue to pay for inputs, namely, other Internet infrastructure.

(p. 9) In the first half of the 1990s, most commercial ISPs tried the same pricing norms governing bulletin boards. For bulletin boards the pricing structure of the majority of services involved a subscription charge (on a monthly or yearly basis) *and* an hourly fee

for usage. For many applications, users could get online for “bursts” of time, which would reduce the size of usage fees.

The emergence of faster and cheaper modems in the mid-1990s and largescale modem banks with lower per-port costs opened the possibility for a different pricing norm, one that did not minimize the time users spent on the telephone communicating with a server. The emergence of low-cost routines for accessing a massive number of phone lines also contributed to a new norm, because it enabled many ISPs to set up modem banks at a scale only rarely seen during the bulletinboard era.

As the Internet commercialized, two broad viewpoints emerged about pricing. One viewpoint paid close attention to user behavior. Users found it challenging to monitor time online, often due to multiple users within one household. The ISPs with sympathy for these user complaints priced their services as unlimited usage for a fixed monthly price. These plans were commonly referred to as *flat-rate* or *unlimited* plans. Another viewpoint regarded user complaints as transitory. Supporters of this view pointed to cellular telephones and bulletin boards as examples where users grew accustomed to pricing-per-minute. The most vocal supporter for this view was one up-and-coming bulletin-board firm, America Online (now known simply as AOL), which had seemed to grow with such usage pricing.

As it turned out, flat rate emerged as the dominant pricing mode for wireline access. There were many reasons for this. For one, the US telephone universal service policy had long subsidized local landline household calls by charging greater prices for business and long-distance calls than for local calls, which were priced at a flat rate per month in almost every state, with the definition of *local* left to a state regulator. Typically *local* was defined over a radius of ten to fifteen miles, sometimes less in dense urban areas, such as Manhattan. Thus, using local rates reduced household user expenses, thereby making the service more attractive. That enabled an ISP to offer service to nontechnical users, betting the users would *not* stay online for long—in effect, anywhere from twenty to thirty hours of time a month at most. Because such light users did not all dial-in on the same day or at the same time of day, the equipment investment did not need to handle all calls at once. With such light users, an ISP could serve a local area with modem bank capacity at anywhere from one-third to one-quarter the size of the total local service base. A large enough user community could thus share equipment, defraying equipment costs for an ISP offering flat rate. Experiments showed that the economies of scale for defraying equipment costs in this way could support approximately 1000 users in a local point of presence. For many dial-up ISPs, this was not a binding constraint.²⁷

Flat rate succeeded for another reason. Other firms entered into the production of web services, complementing what ISPs provided. Flat rate appealed to users who could then choose among a wide array of media and services.²⁸


(p. 10) Throughout 1996, 1997, and 1998, ISPs experimented with hourly limits and penalties for exceeding these caps, yet most users resisted them. The experiments were motivated by many users who stayed online for far more than twenty to thirty hours a month, thereby ruining the key economics that allowed ISPs to share equipment across many users.²⁹ Most such limits were not particularly binding—involving monthly limits ranging from sixty to one hundred hours. Some ISPs tried offering steep discounts for steep limits, such as \$10 discounts for thirty hours a month. Yet few buyers took them, persisting with the slightly more expensive unlimited contracts, typically priced as \$20 per month. In short, despite all these experiments, flat rate dominated transactions during the dial-up era of the commercial Internet.³⁰

5. The Broadband Upgrade

As Internet applications became more popular and more widely adopted, users pushed against the bandwidth limits of dial-up Internet access. That motivated some well-placed firms to deploy and offer broadband access as Internet access. The investment of broadband initiated a second wave of investment in Internet infrastructure after the turn of the millennium. That has been coincident with the presence of more capital deepening in business computing and in many facets of the operations to support it.

The two most common forms were cable modem service and (digital subscriber line) (DSL) service. Cable modem service involved a gradual upgrade to cable plants in many locales, depending on the generation of the cable system.³¹ Broadband over telephone lines involved upgrades to telephone switches and lines to make it feasible to deliver DSL. Both of these choices typically supported higher bandwidth to the household than from the household—thus called asymmetric digital subscriber line (ADSL).

Broadband clearly gave users a better experience than dial-up access.³² Broadband provides households with faster Internet service and thus access to better online applications. Broadband also may allow users to avoid an additional phone line for supporting dial-up. In addition, broadband services are also “always on,” and users perceive that as a more convenient service. It is also generally faster in use. A maximum rate of 14.4K (kilobytes per second) and 28.8K were predominant in the mid-1990s for dial-up modems. The typical bandwidth in the late 1990s was 43K to 51K, with a maximum of 56K. DSL and cable achieved much higher maximum bandwidths, typically somewhere in the neighborhood of a maximum rate of 750K to 3M (megabytes per second), depending on the user choices and vendor configuration. Even higher bandwidth became available to some households later.



[Click to view larger](#)

Figure 1.1 Percentage of Households with Computers and Internet Connections, Selected Years, 1997-2009.


Source: NTIA (2010).

This story is consistent with Figure 1.1, which summarizes US government surveys of broadband use at households. The first survey questions about broadband (p. 11) use appear in 2000 and show a growth in adoption, reaching close to 20 percent of households in 2003, when these surveys

were discontinued for some time.³³ The survey resumed in 2007 and the anticipated trajectory continued, with 50.8 percent of households having broadband in October 2007 and 63.5 percent in October 2009. In the earliest years supply-side issues were the main determinants of Internet availability and, hence, adoption. Cable and telecom operators needed to retrofit existing plants, which constrained availability in many places. In those years, the spread of broadband service was much slower and less evenly distributed than that of dial-up service. Highly populated areas were more profitable due to economies of scale and lower last-mile expenses. As building has removed these constraints, demand-related factors—such as price, bandwidth, and reliability—have played a more significant role in determining the margins between who adopts and who does not.³⁴

Suppliers experimented to find appropriate structures for providing broadband access.³⁵ Countries differed significantly in the extent to which these different delivery channels played a role. Some cable firms built out their facilities to deliver these services in the late 1990s, and many—especially telephone companies—waited until the early to mid-2000s. In some rich countries there was growing use of a third and fourth delivery channel, fiber to the home, and access with mobile modes.³⁶

A similar wave of investment occurred in many developed countries over the first decade of the new millennium. Figure 1.2 shows the subscribers per 100 inhabitants in many countries in 2009.³⁷ Although these numbers must be interpreted with caution, a few facts stand out. A few dozen countries in the Organisation for Economic Co-operation and Development (OECD) experienced substantial adoption of broadband, and many did not. The variance is not surprising. GDP per capita and broadband per capita have a simple correlation of 0.67.³⁸




[Click to view larger](#)

Figure 1.2 OECD Broadband Subscribers per 100 Inhabitants by Technology, June 2009.

Source: OECD.

Figure 1.3 shows the growth of subscribers per 100 inhabitants in the G7, Canada, the United States, United Kingdom, Germany, France, Italy, and Japan, (p. 12) as well as the entire OECD. Though countries differ in the level of broadband use, the similarities among them

are apparent. Adoption of broadband grew in all rich and developed countries.



[Click to view larger](#)

Figure 1.3 Broadband Penetration, G7 Countries.

Source: OECD Broadband Portal, <http://www.oecd.org/sti/ict/broadband>, Table 1i.

To give a sense of worldwide patterns, Table 1.1 presents broadband and dial-up adoption for seven countries—the United States, Canada, United Kingdom, Spain, China, Mexico, and Brazil.³⁹ These seven represent typical experiences in the high-income and middle-income countries of the world. The

broadband data in Table 1.1 come from Point-Topic, a private consultancy.⁴⁰ One fact is immediately obvious. The scale of adoption in the United States and China far outweighs the scale of adoption in any other country. That occurs for two rather obvious economic reasons. The United States and China have much larger populations than the United Kingdom, Spain, and Canada. Although Mexico and Brazil also have (p. 13) large populations, those countries had much lower rates of adoption. In short, the general level of economic development is a major determinant of adoption.

Table 1.1 Broadband Subscribers from Point Topic

Broadband Subscribers from Point Topic in Thousands								
	Year							
Nation	2003	2004	2005	2006	2007	2008	2009	CAGR
Brazil	634	1,442	2,671	4,278	5,691	7,509	9,480	47.2%
Canada	3,706	4,829	5,809	6,982	8,001	8,860	9,528	14.4%
China	-	11,385	20,367	30,033	41,778	54,322	68,964	35.0%
Mexico	234	429	1,060	1,945	3,106	4,774	7,836	65.1%
Spain	1,401	2,524	3,444	5,469	7,322	8,296	9,023	30.5%
United Kingdom	960	3,734	7,203	10,983	13,968	16,282	17,641	51.6%
United States	16,042	28,770	37,576	47,489	58,791	67,536	77,334	25.2%

Source: Greenstein and McDevitt (2011)

The popularity of wireline broadband access has reopened questions about pricing. Following precedent, most broadband providers initially offered unlimited plans. High volume applications—such as video downloading and peer-to-peer applications—placed pressures on capacity, however. This motivated providers to reconsider their flat-rate pricing plans and impose capacity limitations. Hence, a decade and half after the blossoming of the commercial Internet there was considerable speculation about whether video applications would generate a new norm for pricing access. The answer partly depended on the regulatory regime governing Internet access and remains open as of this writing.⁴¹

6. Expanding Modes of Access

In the decade after the millennium another and new set of mobile broadband services began to gain market traction with businesses and households. The first were providers known as wireless ISPs, or WISPs for short. They provided access via a variety of technologies, such as satellite, high-speed WiFi (wireless local area network that uses high-frequency radio signals to transmit and receive data), WiMax (stands for worldwide interoperability for microwave access), and other terrestrial fixed-point wireless delivery modes. These providers primarily served low-density locations where the costs of wireline access were prohibitive.⁴²

Another wireless mode became popular, smart phones. Though smart phones had been available in a variety of models for many years, with Blackberry being among the most popular, it is commonly acknowledged that the category began (p. 14) to gain in adoption after the introduction of the Apple iPhone in 2007. As of this writing, reports suggest the Apple iPhone, Google Android, and new Blackberry designs dominate this product category for the time being. Yet competitive events remain in flux. Competitive responses organized by Palm, Microsoft, and Nokia have been attempted, and those firms and others will continue to attempt more. In short, irrespective of which firms dominate, smart phones are an important new access point.

The economics of smart-phone use remain in flux as well. It is unclear whether the majority of users treat their smart phones as substitutes to their home broadband use. Smart phones provide additional mobility, and that might be a valuable trait by itself. If smart phones are simply additional services due to mobility, then the economic value of smart phone use should be interpreted one way. If the additional services are partial or complete substitutes, then the economic value of smart phone use should be interpreted another.

Another open question concerns the boundaries between home and business use. With wireline broadband this is less of a question because the destination of the location for Internet access largely identifies its buyer (home or business). Smart phones, however, sell both to home and business, thus obliterating what had been a bright line between these kinds of customers and altering the economic analysis of the adoption decision and its impact. This also alters the potential revenue streams for Internet access, as well as the geographic distribution of use, which could lead to changes in the funding for more Internet infrastructure, as well as the geographic features of investment.

In addition, because of leapfrogging of mobile over fixed broadband in emerging economies, mobile broadband may be the first broadband experience for many people. So not only is it unclear whether mobile broadband substitutes or complements fixed broadband, but the extent of substitutability could vary substantially by country according to each country's stage of infrastructure development.

The changing modes of access opened multiple questions about changing norms for pricing-data services. Wireless data services over cellular networks generally did not use flat-rate pricing and often came with explicit charges for exceeding capacity limits, for example. As of this writing, the norm for these plans was, and continues to be, an open question.

7. Strategic Behavior from Suppliers

When it first deployed, the Internet was called a “network of networks.” Although the phrase once had meaning as a description of the underlying configuration of infrastructure, it is misleading today. Leading firms and their business partners view the commercial Internet through the same lens they view the rest of computing. To them, the Internet is a market in which to apply their platform (p. 15) strategies. In short, the commercial Internet should be called a “network of platforms.”

The list of important platforms in Internet infrastructure is long, and many involve familiar firms, such as Google, Apple, Cisco, Microsoft, and Intel. As a result, the platform strategies of private firms shape the evolution of Internet infrastructure.⁴³ This was one of the biggest changes wrought by introducing private firms into the supply of Internet infrastructure.

Commercial firms regard platforms as reconfigurable clusters and bundles of technical standards for supporting increasing functionality.⁴⁴ From a user perspective, platforms

usually look like “standard bundles” of components, namely, a typical arrangement of components for achieving functionality.⁴⁵

Platforms matter because platform leaders compete with others. That usually leads to lower prices, competition in functionality, and accelerated rollout of new services. Many platform leaders also develop new functionality and attach it to an existing platform.⁴⁶ Generally, that brings new functionality to users. A platform leader also might seek to fund businesses outside its area of expertise, if doing so increases demand for the core platform product.⁴⁷ Hence, in the commercial Internet, the strategic priorities of its largest providers tend to shape the innovative characteristics of Internet infrastructure.

Today many observers believe that Google—which did not even exist when the Internet first commercialized in the mid-1990s—has the most effective platform on the Internet. Hence, its behavior has enormous economic consequences. For example, sometimes it makes code accessible to programmers for mash-ups—for example, building services that attract developers and users with no direct way to generate revenue.⁴⁸ Sometimes its investments have direct effects on other Internet infrastructure firms. For example, Google operates an enormous global Internet backbone, as well as CDN and caching services, for its own operations, making it one of the largest data-transit providers in the world today. The firm also uses its platform to shape the actions of others. Many other firms expend considerable resources optimizing their web pages to appear high on Google's search results, and Google encourages this in various ways. For example, it lets potential advertisers, who will bid in Google's auction, know which words are the most popular.

Networking equipment provider Cisco is another prominent platform provider for Internet infrastructure, having grown to become a large provider of equipment for business establishments. For many years, Cisco made most of its profit from selling hubs and routers, so the platform strategy was rather straightforward. Cisco aspired to developing closely related businesses, offering users a nearly integrated solution to many networking problems. At the same time, Cisco kept out of service markets and server applications, leaving that to integrators, consultants, and software vendors. That way, Cisco did not compete with its biggest business partners. More recently, however, Cisco branched into consumer markets (with its purchase of Linksys). The firm also has moved into some server (competing with HP) and some software/service areas related to videoconferencing and telepresence (p. 16) (by purchasing Webex, for example). Cisco no longer draws the boundary where it used to, and it is unclear how wide a scope the firm wants its platform to cover.

Microsoft is perhaps the next best known platform provider whose business shapes Internet infrastructure. In the early 1990s, Microsoft offered TCP/IP compatibility in

Windows as a means of enhancing its networking software, as well as to support functionality in some of its applications, such as Exchange. In the mid-1990s, Microsoft offered a browser, partly as a gateway toward developing a broader array of web services, and partly for defensive purposes, to establish its proprietary standards as dominant.⁴⁹ Although Microsoft continues to support these commercial positions and profit from them, the firm has not had as much success in other aspects of its commercial Internet ventures. MSN, search, mobile OS, and related activities have not yielded sustained enviable success (yet). Only the firm's investments in Xbox Live have generated a significant amount of Internet traffic, and it continues to push the boundaries of large-scale multiplayer Internet applications.

Another PC firm, Intel, has an Internet platform strategy. Intel's most important Internet activity came from sponsoring a Wi-Fi standard for laptops under the Centrino brand in 2003.⁵⁰ To be clear, this *did not* involve redesigning the Intel microprocessor, the component for which Intel is best known. It did, however, involve redesigning the motherboard for desktop PCs and notebooks by adding new parts. This redesign came with one obvious benefit: It eliminated the need for an external card for the notebook, usually supplied by a firm other than Intel and installed by users (or OEMs—original equipment manufacturers) in an expansion slot. Intel also hoped that its endorsement would increase demand for wireless capabilities within notebooks using Intel microprocessors by, among other things, reducing their weight and size while offering users simplicity and technical assurances in a standardized function. Intel hoped for additional benefits for users, such as more reliability, fewer set-up difficulties, and less frequent incompatibility in new settings. Intel has helped fund conformance-testing organizations, infrastructure development, and a whole range of efforts in wireless technology. More recently, it has invested heavily in designing and supporting other advanced wireless standards, such as WiMax.

As with many other aspects of commercialization, the importance of platforms is both cause for celebration and a source of concern. It is positive when platform strategies help firms coordinate new services for users. Why is the emergence of platform providers a concern? In short, in this market structure the private incentives of large dominant firms determine the priorities for investment in Internet infrastructure. Under some circumstances dominant platform firms have incentives to deny interconnection to others, to block the expansion of others, and, in an extreme case, to smother the expansion of new functionality by others. When Internet alarmists worry about conduct of proprietary platforms, they most fear deliberate introduction of incompatibilities between platforms, and other conduct to deliberately violate end-to-end principles.⁵¹ Microsoft, AOL, Intel, Comcast, and WorldCom have all shown tendencies toward such behavior in specific

episodes. (p. 17) Open-source advocates who claim that they benefit the Internet usually mean they are preventing defensive activity by leaders with defensive tendencies.

More recently, a range of behavior from Apple, Facebook, Twitter, American Airlines, and the *Wall Street Journal* have raised fears about the “splintering” of the Internet.⁵²

Splintering arises when users lose the ability to seamlessly move from one application to another or vendors lose a common set of platforms on which to develop their applications. Steve Jobs’ decision not to support Flash on the iPhone and iPad is one such recent example. So is Twitter's resistance to having its services searched without payment from the search engines, and Facebook has taken a similar stance. American Airlines refused to let Orbitz see its prices without modifying its website to accommodate new services American Airlines wanted to sell to frequent fliers, and Orbitz refused. As a result, the two companies no longer cooperate. For some time the *Wall Street Journal* has refused to let users search extensively in its archives without a subscription, and its management openly discusses aspirations to gain a fee from search engines, much as Facebook and Twitter did.

What economic incentives lie behind these concerns? In a nutshell the answer begins thusly: no law compels any firm to interconnect with any other. Every firm always has the option to opt out of the Internet, and/or do something slightly less dramatic, such as opt out of commonly used standards, and/or try to get business partners to use proprietary standards. At any given moment, that means a firm in a strong competitive position will have incentives to raise switching costs to its installed base of developers and users, and/or deter migration of users and developers to competing platforms. A variety of strategic options might contribute to such goals, such as designing proprietary standards or blocking others from using a rival firm's standard designs.

Concerns about splintering also arise because suppliers must cooperate in order to deliver services. Because one party's cost is another party's revenue, firms have incentives to sacrifice aspects of cooperation in the attempt to gain rents. Such behavior is not necessarily in users’ interests.

Consider the negotiations between Cogent and Sprint, for example, which broke down after a peering dispute.⁵³ Cogent refused to pay Sprint after Sprint insisted Cogent had not met their obligations under a peering agreement. After a long stand-off, Sprint's management decided to shut down its side of the peering. That action had consequences for users on both networks who did not multihomed, that is, did not use more than one backbone firm. One set of exclusive Sprint users could not reach another set of exclusive Cogent users.⁵⁴ To make a long story short, users of both carriers were angry, and Sprint's management gave in after a few days. Soon after, the two firms came to a long-term agreement whose details were not disclosed publicly. Notice how inherently

awkward the negotiations were: Cogent necessarily interacted or exchanged traffic with the very firm with which it competes, Sprint.⁵⁵

Another set of cases illustrates how the interests of one participant may or may not intersect with the interests of all participants. For example, consider Comcast's unilateral declaration to throttle peer-to-peer (P2P) applications on its (p. 18) lines with resets.⁵⁶

This case contains two economic lessons. On the one hand, it partially favors giving discretion to Comcast's management. Management could internalize the externality one user imposes on others—managing traffic for many users' general benefit. That is, P2P applications, like Bit-Torrent, can impose large negative externalities on other users, particularly in cable architectures during peak-load time periods. Such externalities can degrade the quality of service to the majority of users without some sort of limitation or restriction. On the other hand, Comcast's behavior shapes at least one additional provider of applications, future entrepreneurs, many of whom are not present. It would be quite difficult for Comcast and future entrants to reach a bargain because some of them do not even exist yet. Eventually the FCC intervened with Comcast, issuing an order to cease blocking, which led to a court case over its authority to issue such an order. As of this writing, the full ramifications of this court case have not played themselves out.

Another case involving Comcast also illustrates the open-ended nature of the cooperation between firms. In November 2010, Comcast, the largest provider of broadband Internet access in the United States, entered into a peering dispute with Level 3, one of the backbone firms with which it interconnected. Level 3 had made an arrangement with Netflix, a video provider, and this had resulted in Level 3 providing more data to Comcast than Comcast provided to Level 3. Comcast demanded that Level 3 pay for giving more data to Comcast than Comcast gave in return, to which Level 3 agreed.⁵⁷ This agreement was significant because it was the first time an ISP did not pay the backbone provider for transit services, but instead, the provider paid a large ISP to reach end users. As of this writing, it is an open question how common such agreements will become, and whether they will alter the general flow of dollars among Internet infrastructure firms.

8. Governance for Internet Infrastructure

A number of organizations play important roles in supporting the design, upgrading, and operations of Internet infrastructure. The most notable feature of the governance of Internet infrastructure are the differences with the governance found in any other communications market.

One notable feature of this structure was the absence of much government directive or mandate. It is incorrect to say that the government was uninvolved: After all, the NSF and Department of Defense both had played a crucial role in starting and sponsoring organizations that managed and improved the operations of the Internet.⁵⁸ Rather, the commercial Internet embodied the accumulation of multiple improvements suggested through a process of consensus in committees, and that consensus depended in large part on private action, what economists call (p. 19) “private orderings.”⁵⁹ Unlike any other communication network, governments did not play a substantial role in these private orderings.

The organization that governs the upgrades to TCP/IP is the Internet Engineering Task Force (IETF). It was established prior to the Internet's privatization, and continued as a nonprofit organization after its commercialization. Today it hosts meetings that lead to designs that shape the operations of every piece of equipment using TCP/IP standards.⁶⁰ Many of these decisions ensured that all complying components would interoperate. Today decisions at the IETF have enormous consequences for the proprietary interests of firms.

Standards committees had always played some role in the computer market, and they played a similar role in the shaping of Internet infrastructure. The Institute of Electrical and Electronics Engineers (IEEE), for example, made designs that shaped the LAN market, modem, and wireless data communications markets.⁶¹

Aside from the absence of government mandates, these groups also were notable for the absence of dominant firms. They were not beholden to the managerial auspices of AT&T or IBM, or any other large firm, for example. Though all those firms sent representatives who had a voice in shaping outcomes, these institutions were characterized by divided technical leadership.

That does not imply that all these organizations conducted their business in a similar manner. On the contrary, these forums differed substantially in their conduct.⁶² The World Wide Web Consortium (W3C) offers an illuminating comparison. Berners-Lee forecast the need for an organization to assemble and standardize pieces of codes into a broad system of norms for operating in the hyper-text world. He founded the World Wide Web Consortium for this purpose. In 1994 he established the offices for the W3C in Cambridge, Massachusetts, just in time to support an explosion of web-based services that took advantage of existing Internet infrastructure.

Berners-Lee stated that he had wanted a standardization process that worked more rapidly than the IETF but otherwise shared many of its features, such as full documentation and unrestricted use of protocols. In contrast to the IETF, the W3C would

not be a bottom-up organization with independent initiatives, nor would it have unrestricted participation. Berners-Lee would act in a capacity to initiate and coordinate activities. To afford some of these, his consortium would charge companies for participating in efforts and for the right to keep up-to-date on developments.⁶³

The governance structures for the IETF and the W3C also can be compared to what it is not, namely, the next closest alternative for global networking—the Open Systems Interconnection model, a.k.a., OSI seven-layer model. The OSI was a formal standard design for interconnecting networks that arose from an international standards body, reflecting the representation of multiple countries and participants. The processes were quite formal. The network engineering community in the United States preferred their bottom-up approach to the OSI top-down approach and, when given the opportunity, invested according to their preferences.⁶⁴

The lack of government involvement could also be seen in other aspects of the Internet in the United States. For example, the Federal Communications Commission (FCC) refrained from mandating most Internet equipment design decisions. Just as the FCC had not mandated Ethernet design standards, so it let the spectrum become available for experiments by multiple groups who competed for *wireless* Ethernet standards, which eventually became Wi-Fi. Similarly, the FCC did not mandate a standard for modems other than to impose requirements that limited interference. It also did not mandate an interconnection regulatory regime for Internet carriers in the 1990s.⁶⁵

The US government's most visible involvement in governance has come with its decisions for the Internet Corporation for Assigned Numbers and Names (ICANN), an organization for governing the allocation of domain names. The NSF originally took responsibility for domain names away from the academic community prior to privatizing the Internet, giving it to one private firm. The Department of Commerce established a nonprofit organization to provide oversight of this firm, with the understanding that after a decade ICANN would eventually become a part the United Nations.⁶⁶ This latter transfer never took place, and, as of this writing, ICANN remains a US-based nonprofit corporation under a charter from the Commerce Department.

One other notable innovative feature of Internet infrastructure is its reliance on the behavioral norms and outcomes of open-source projects. This had substantial economic consequences, establishing behavior norms for information transparency that had never before governed the design of equipment for a major communication market. Key aspects of Internet infrastructure embedded designs that emerged from designs that any firm or user could access without restriction, and to which almost any industry participant could make contributions.

One well-known open-source project was Linux, a basis for computer operating systems. It was begun by Linus Torvald in the early 1990s as a derivative, or “fix,” to Unix. It was freely distributed, with alternating releases of a “beta” and “final” version. Starting around 1994 to 1995, about the same time as the commercialization of the Internet, Linux began to become quite popular. What had started as a small project caught the attention of many Unix users, who started contributing back to the effort. Many users began converting from proprietary versions of Unix (often sold by the hardware manufacturers) and began basing their operating systems on Linux, which was not proprietary. This movement gained so much momentum that Linux-based systems became the most common server software, especially for Internet servers.⁶⁷

Apache was another early project founded to support and create “fixes” for the HTTP web server originally written by programmers at the National Center for Super Computing Applications (NCSA) at the University of Illinois. By 2006, more than 65 percent of websites in the world were powered by the Apache HTTP web server.⁶⁸ Apache differed from many other open-source organizations in that contributors “earned” the right to access the code. To be a contributor one had to be working on at least one of Apache's projects. By 2006, the organization had an active and large contributor base.

(p. 21) Perhaps the most well-known open source format was the least technical. It originated from something called a *wiki*. Developed in 1995 by Ward Cunningham, a software engineer from Portland, Oregon, wikis can either be used in a closed work group or used by everyone on the Internet. They originally were formed to replicate or make a variation on existing products or services, with the purpose of fixing bugs within the various systems. Accordingly, wikis were first developed and intended for software development but had grown out of that first use and became applied to a multitude of applications. In short, wikis became the essential software infrastructure upon which virtually all major Internet media applications are built.

A particular popular application of wikis, Wikipedia, garnered worldwide attention. In the case of Wikipedia, the format was applied to the development of textual and nontextual content displayed on the Web. It is an online-only encyclopedia. The content is user-created and edited. As its homepage proudly states, it is “The Free Encyclopedia That Anyone Can Edit.” The site has always been free of charge and never accepted advertising.⁶⁹ Wikipedia beat out Microsoft's Encarta for the honor of the Internet's top research site in 2005, a position that it has held ever since.⁷⁰

An experimental form of copyright, the creative commons license, spread extraordinarily fast. This license, founded in only 2001, is used by over 30 million websites today.⁷¹ It has begun to play a prominent role in online experimentation and everyday activity. Creative commons licenses help organizations accumulate information in a wide array of new

business formats. Flickr is one successful example, having recently passed the milestone of four billion photos on its site. The creative commons license also is employed by numerous Web2.0 initiatives and new media, which support online advertising.

After so much experience with open source it is no surprise that the major participants in Internet infrastructure no longer leave these institutions alone. The standardization organizations find their committees filled with many interested participants, some with explicit commercial motives and some not. These institutions show signs of the stress, chiefly in a slowing down in their decision making, if they reach decisions at all. Perhaps that should also be cause for celebration, since it is an inevitable symptom of commercial success and the large commercial stakes for suppliers.⁷²

9. Broadband Policy

Many governments today, especially outside the United States, are considering making large subsidies for broadband investments. Some governments, such as those of South Korea and Australia, have already done so, making next-generation broadband widely available. Many years of debate in the United States led to the emergence of a National Broadband Plan, released in March 2010.⁷³ Related debates also led to a large European framework for the governance of investments in broadband.⁷⁴

(p. 22) Some of the key issues can be illustrated by events in the United States. At the outset of the commercial Internet, policy favored allowing firms to invest as they please. During the latter part of the 1990s, policy did not restrict the ability of firms to respond to exuberant and impatient demand for new Internet services. After a time, the infrastructure to support those new services became far larger than the demand for services.⁷⁵ After the dot-com boom came to a bust, the United States found itself with excessive capacity in backbone and many other infrastructure facilities to support Internet services.

In the decade after, policy continued to favor privately financed investment. That resulted in a gradual building of broadband. Why did private supply favor gradualism? In short, aside from perceptions of overcapacity, few executives at infrastructure firms would ever have deliberately invested resources in an opportunity that was unlikely to generate revenue until much later, especially ten to twenty years later. Corporate boards would not have approved of it, and neither would stockholders. One of the few firms to attempt such a strategy was Verizon, which unveiled a program to build fiber to the home in the latter part of the first decade after the millennium. Due to low take-up, Verizon did not fully build these services in all its territory.⁷⁶

Most arguments for building next-generation Internet broadband ahead of demand faced large political obstacles. Consider one justification, economic experimentation, namely, better broadband clearly helps experimentation in applications by making users better customers for online ads and electronic retailing. Although the monetary costs are easy to tally, the benefits are not. Relatedly, the costs are focused, but the gains are diffuse, thus making it difficult to show that broadband caused the associated gain, even if, broadly speaking, everyone recognizes that broadband raised firms' productivity and enhanced users' experience. Accordingly, financing broadband would involve a general tax on Yahoo and Amazon and Google and YouTube and other national electronic retailers and application providers who benefit from better broadband. Needless to say, considerable political challenges interfere with the emergence of such schemes. Some countries, such as Korea, have managed to come to such a political agreement, but these examples are the exception, not the rule.

Another set of policies considers subsidizing rural broadband, that is, subsidizing the costs of building wireline broadband supply in high-cost areas.⁷⁷ Since private supply already covers the least costly areas to supply, only a small fraction of potential households benefit from such subsidies, and the costs of subsidizing buildouts are especially high. Such subsidies face numerous challenges. The US National Broadband Plan provides an excellent summary of the issues. Many of the justifications for these subsidies are noneconomic in nature—aimed at increasing civic engagement among rural populations, increasing obtainment of educational goals among children, or increasing the likelihood of obtaining specific health benefits. Accordingly, the decision to provide such subsidies is often a political decision rather than purely an economic one.

Another open question concerns the governance of deployment of access networks. In Europe the governments have chosen a structure that essentially (p. 23) separates ownership of transmission facilities from ownership of content, and mandates interconnection for many rivals at key points.⁷⁸ In the United States there was a preference for private provision of backbone and access networks in the first decade of the millennium, and a light-handed degree of regulatory intervention, so providers were not required to offer interconnection to others.

No simple statement could characterize the changing norms for regulating Internet infrastructure. For example, as of this writing, at the federal level, there are initiatives under way to adopt formal policies for ensuring the openness of Internet access.⁷⁹ At the local level, there are a range of initiatives by municipalities to provide local access in competition with private suppliers.⁸⁰ There is also effort to limit municipal intervention in the provision of access or deter state limitations on local initiatives.⁸¹

10. Summary

No single administrative agency could possibly have built and managed the commercial network that emerged after the privatization of the Internet. The shape, speed, growth, and use of the commercial Internet after 1995 exceeded the ability of any forecaster inside or outside government circles. The value chain for Internet services underwent many changes after the Internet was privatized. More investment from private firms, and more entry from a range of overlays and new applications, altered nearly every aspect of the structure of the value chain. This evolution occurred without explicit directives from government actors, with only a light hand of directives, and with astonishing speed.

Many key events in Internet infrastructure took place within the United States in the first decade and a half of the commercial Internet, but this appears to be less likely as the commercial Internet grows large and more widespread. While the United States continues to be the source of the largest number of users of Internet services, and the single greatest origin and destination for data traffic, the US position in the global Internet value chain will not—indeed, cannot—remain dominant. That should have enormous consequences for the evolution of the structure of global Internet infrastructure because many countries insist on building their infrastructure according to principles that differ from those that governed the first and second waves of investment in the United States. The boundaries between public and private infrastructure should change as a result, as should the characteristics of the governance and pricing of Internet infrastructure.

It is no surprise, therefore, that many fruitful avenues for further economic research remain open. For example, what frameworks appropriately measure the rate of return in investment in digital infrastructure by public and private organizations? Through what mechanisms does advance Internet infrastructure produce (p. 24) economic growth, and to which industries in which locations do most of the positive and negative effects flow? What factors shape the effectiveness of different governance structures for open structures, such as those used by the IETF? What is the quantitative value of these novel governance structures?

For the time being there appears to be no cessation in the never-ending nature of investment in Internet infrastructure. Indeed, as of this writing, many questions remain open about the value of different aspects of IT in the long run, and firms continue to explore approaches to creating value. Virtually all participants in these markets expect continual change, as well as its twin, the absence of economic tranquility.

References

Abbate, J., 1999. *Inventing the Internet*, MIT Press: Cambridge, Mass.

Aizcorbe, A., K. Flamm, and A. Khursid, 2007. "The Role of Semiconductor Inputs in IT Hardware Price Decline: Computers vs. Communications," in (eds.) E. R. Berndt and C. M. Hulten, *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, University of Chicago Press, pp. 351–382.

Anderson, C., and M. Wolff, 2010. "The Web Is Dead. Long Live the Internet," *Wired*, September. http://www.wired.com/magazine/2010/08/ff_webrip/.

Arora, A., and F. Bokhari, 2007. "Open versus Closed Firms and the Dynamics of Industry Evolution," *Journal of Industrial Economics*, 55(3), 499–527.

Augereau, A., S. Greenstein, and M. Rysman, 2006. "Coordination versus Differentiation in a Standards War: 56K Modems," *Rand Journal of Economics*, 34 (4), 889–911.

Baumol, W., and twenty six economists. 2006. Economists' Statement on U.S. Broadband Policy (March 2006). AEI-Brookings Joint Center Working Paper No. 06-06-01. Available at SSRN: <http://ssrn.com/abstract=892009>.

Berners-Lee, T., and M. Fischetti, 1999. *Weaving the Web, The Original Design and Ultimate Destiny of the World Wide Web*, Harper Collins, New York.

Besen, S., P. Milgrom, B. Mitchell, and P. Sringamesh, 2001. "Advances in Routing Technologies and Internet Peering Agreements," *American Economic Review*, May, pp. 292–296.

Blumenthal, M. S., and D. D. Clark, 2001. "Rethinking the Design of the Internet: The End-to-End Arguments vs. The Brave New World." In (eds.) B. Compaine and S. Greenstein, *Communications Policy in Transition: The Internet and Beyond*. MIT Press: Cambridge, Mass., pp. 91–139.

Bresnahan, T., 1999. "The Changing Structure of Innovation in Computing" in (ed.) J. A. Eisenach and T. M. Lenard, *Competition, Convergence and the Microsoft Monopoly: Antitrust in the Digital Marketplace*, Kluwer Academic Publishers, Boston, pp. 159–208.

(p. 29)

Bresnahan, T., and Pai-Ling Yin, 2007. "Standard Setting in Markets: The Browser Wars," in (eds.) S. Greenstein and V. Stango, *Standards and Public Policy*, Cambridge University Press; Cambridge, UK. pp. 18–59.

Bresnahan, T., and S. Greenstein, 1999. "Technological Competition and the Structure of the Computer Industry," *Journal of Industrial Economics*, March, pp. 1-40.

Bresnahan, T., S. Greenstein, and R. Henderson, 2011. "Schumpeterian Competition and Diseconomies of Scope: Illustrations from the Histories of IBM and Microsoft." Forthcoming in (eds.) J. Lerner and S. Stern, *The Rate and Direction of Inventive Activity, 50th Anniversary*, National Bureau of Economic Research. <http://www.nber.org/books/lern11-1>, accessed March, 2012.

Burgelman, R., 2007. "Intel Centrino in 2007: A New Platform Strategy for Growth." Stanford University Graduate School of Business, SM-156.

Chiao, B., J. Tirole, and J. Lerner, 2007. "The Rules of Standard Setting Organizations," *Rand Journal of Economics*, 34(8), 905-930.

Clark, D., W. Lehr, S. Bauer, P. Faratin, R. Sami, and J. Wroclawski, 2006. "Overlays and the Future of the Internet," *Communications and Strategies*, 63, 3rd quarter, pp. 1-21.

Crandall, R., 2005. "Broadband Communications," in (eds.) M. Cave, S. Majumdar, and Vogelsang, *Handbook of Telecommunications Economics*, pp. 156-187. Amsterdam, The Netherlands: Elsevier.

Cusumano, M., and D. Yoffie, 2000. *Competing on Internet Time: Lessons from Netscape and Its Battle with Microsoft*. New York: Free Press.

Dalle, J.-M., P. A. David, R. A. Ghosh, and F. Wolak, 2004. "Free & Open Source Software Developers and 'the Economy of Regard': Participation and Code-Signing in the Modules of the Linu Kernel," Working paper, SIEPR, Stanford University, Open Source Software Project http://siepr.stanford.edu/programs/OpenSoftware_David/NSFOSF_Publications.html

Dedrick, J., and J. West, 2001. "Open Source Standardization: The Rise of Linu in the Network Era," *Knowledge, Technology and Policy*, 14 (2), 88-112.

Doms, M., and C. Forman, 2005. "Prices for Local Area Network Equipment," *Information Economics and Policy*, 17(3), 365-388.

Downes, T., and S. Greenstein, 2002. "Universal Access and Local Internet Markets in the U.S.," *Research Policy*, 31, 1035-1052.

Downes, T., and S. Greenstein, 2007. "Understanding Why Universal Service Obligations May Be Unnecessary: The Private Development of Local Internet Access Markets." *Journal of Urban Economics*, 62, 2-26.

Evans, D., A. Hagiu, and R. Schmalensee (2006). *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries*, MIT Press; Cambridge, Mass.

Federal Communications Commission, 2010a. *National Broadband Plan, Connecting America*, <http://www.broadband.gov/>.

Federal Communications Commission, 2010b. "In the Matter of Preserving the Open Internet Broadband Industry Practices," GN Docket No. 09-191, WC Docket No. 07-52, December 23, 2010. <http://www.fcc.gov/>.

Forman, C., and A. Goldfarb, 2006. "Diffusion of Information and Communications Technology to Business," in (ed.) T. Hendershott, *Economics and Information Systems, Volume 1*, Elsevier, pp. 1-43.

Forman, C., A. Goldfarb, and S. Greenstein, 2005. "How did Location Affect Adoption of the Internet by Commercial Establishments? Urban Density versus Global Village," *Journal of Urban Economics*, 58(3), 389-420. (p. 30)

Fosfuri, A., M. Giarratana, and A. Luzzi, 2005. "Firm Assets and Investments in Open Source Software Products." Druid Working Paper No. 05-10. Copenhagen Business School.

Friedan, R., 2001. "The Potential for Scrutiny of Internet Peering Policies in Multilateral Forums," in (eds.) B. Compaine and S. Greenstein, *Communications Policy in Transition, The Internet and Beyond*, MIT Press; Cambridge, Mass., 159-194.

Gawer, A., 2009, *Platforms, Innovation and Competition*, Northampton, Mass.: Edward Elgar.

Gawer, A., and M. Cusumano, 2002. *Platform Leadership: How Intel, Microsoft and Cisco Drive Innovation*. Boston, Mass.: Harvard Business School Press.

Gawer, A., and R. Henderson, 2007. "Platform Owner Entry and Innovation in Complementary Markets: Evidence from Intel." *Journal of Economics and Management Strategy*, Volume 16 (1), 1-34.

Gilles, J., and R. Cailliau, 2000, *How the Web Was Born*, Oxford, UK: Oxford University Press.

Goldfarb, A., 2004. "Concentration in Advertising-Supported Online Markets: An Empirical Approach." *Economics of Innovation and New Technology*, 13(6), 581-594.

Goldstein, F., 2005. *The Great Telecom Meltdown*. Boston: Artech House.

Greenstein, S., 2006. "Wikipedia in the Spotlight." Kellogg School of Management, case 5-306-507, http://www.kellogg.northwestern.edu/faculty/kellogg_case_collection.aspx.

Greenstein, S., 2007a. "Economic Experiments and Neutrality in Internet Access Markets," in (eds.) A. Jaffe, J. Lerner and S. Stern, *Innovation, Policy and the Economy*, Volume 8. Cambridge, Mass.: MIT Press, pp. 59-109.

Greenstein, S., 2007b. "The Evolution of Market Structure for Internet Access in the United States," in (eds.) W. Aspray and P. Ceruzzi, *The Commercialization of the Internet and its Impact on American Business*. Cambridge, Mass.: MIT Press, pp. 47-104.

Greenstein, S., 2009a. "Open Platform Development and the Commercial Internet." In (ed.) A. Gawer, *Platforms, Innovation and Competition*, Northampton, Mass.: Edward Elgar, pp. 219-250.

Greenstein, S., 2009b. "Glimmers and Signs of Innovative Health in the Commercial Internet," *Journal of Telecommunication and High Technology Law*, pp. 25-78.

Greenstein, S., 2010. "The Emergence of the Internet: Collective Invention and Wild Ducks." *Industrial and Corporate Change*. 19(5), 1521-1562.

Greenstein, S., 2011. "Nurturing the Accumulation of Innovations: Lessons from the Internet," in (eds.) Rebecca Henderson and Richard Newell, *Accelerating Innovation in Energy: Insights from Multiple Sectors*, University of Chicago Press, pp. 189-224.

Greenstein, S., and R. McDevitt. 2009. "The Broadband Bonus: Accounting for Broadband Internet's Impact on U.S. GDP." NBER Working paper 14758. <http://www.nber.org/papers/w14758>.

Greenstein, S., and R. McDevitt, 2011. "Broadband Internet's Impact on Seven Countries," in (ed.) Randy Weiss, *ICT and Performance: Towards Comprehensive Measurement and Analysis*. Madrid: Fundacion Telefonica, pp. 35-52.

Greenstein, S., and J. Prince, 2007. "The Diffusion of the Internet and the Geography of the Digital Divide," in (eds.) R. Mansell, C. Avgerou, D. Quah, and R. Silverstone, *Oxford Handbook on ICTs*, Oxford University Press, pp. 168-195. (p. 31)

Haigh, T., 2007. "Building the Web's Missing Links: Portals and Search Engines" in (eds.) W. Aspray and P. Ceruzzi, *The Internet and American Business*, MIT Press.

Kahin, B., and B. McConnell, 1997. "Towards a Public Metanetwork; Interconnection, Leveraging and Privatization of Government-Funded Networks in the United States," in (eds.) E. Noam and A. Nishuilleabhain, *Private Networks Public Objectives*, Elsevier, Amsterdam, the Netherlands. pp. 307–321.

Kahn, R., 1995. "The Role of Government in the Evolution of the Internet," in (ed.) National Academy of Engineering, *Revolution in the U.S. Information Infrastructure*. Washington, D.C.: National Academy Press, 13–24.

Kende, M., 2000. "The Digital Handshake: Connecting Internet Backbones," Working Paper No. 32, Federal Communications Commission, Office of Planning and Policy, Washington, D.C.

Kesan, J. P., and R. C. Shah, 2001. "Fool Us Once, Shame on You—Fool Us Twice, Shame on Us: What We Can Learn from the Privatizations of the Internet Backbone Network and the Domain Name System," *Washington University Law Quarterly*, 79, 89–220.

Laffont, J.-J., J. S. Marcus, P. Rey, and J. Tirole, 2001. "Internet Peering," *American Economic Review, Papers and Proceedings*, 91, . 287–292.

Laffont, J.-J., J. S. Marcus, P. Rey, and J. Tirole, 2003. "Internet Interconnection and the Off-net Pricing Principle," *Rand Journal of Economics*, 34(2), 370–390.

Leiner, B., V. Cerf, D. Clark, R. Kahn, L. Kleinrock, D. Lynch, J. Postel, L. Roberts, and S. Wolff, 2003. *A Brief History of the Internet, Version 3.32*, Last revised, 10 December, 2003. <http://www.isoc.org/internet/history/brief.shtml>, downloaded August, 2009.

Lerner, J., and J. Tirole, 2002. "Some Simple Economics of Open Source Software," *Journal of Industrial Economics*, 50 (2), 197–234.

Lessig, L., 1999, *Code and Other Laws of Cyber Space*, New York: Basic Books.

Mackie-Mason, J., and J. Netz, 2007., "Manipulating Interface Standards as an Anticompetitive Strategy," in (eds.) Shane Greenstein and Victor Stango, *Standards and Public Policy*, Cambridge Press; Cambridge, Mass., pp. 231–259.

Marcus, J. S., 2008. "IP-Based NGNs and Interconnection: The Debate in Europe," *Communications & Strategies*, No. 72, p. 17.

Marcus, J. S., and D. Elixmann, 2008. "Regulatory Approaches to Next Generation Networks (NGNs): An International Comparison." *Communications and Strategies*, 69 (1st quarter), 1–28.

Mockus, A., R. T. Fielding, and J. D. Herbsleb, 2005. "Two Case Studies of Open Source Software Development: Apache and Mozilla," in (ed.) J. Feller, *Perspectives on Free and Open Source Software*. Cambridge, MA: MIT Press; Cambridge, pp 163-210.

Mowery, D. C., and T. S. Simcoe, 2002a. "The Origins and Evolution of the Internet," in R. Nelson, B. Steil, and D. Victor (eds.), *Technological Innovation and Economic Performance*. Princeton, N.J.: Princeton University Press, 229-264.

Mowery, D. C. and T. S. Simcoe, 2002b. "Is the Internet a U.S. Invention? An Economic and Technological History of Computer Networking." *Research Policy*, 31(8-9), 1369-1387.

Mueller, M.L., 2002. *Ruling the Root: Internet Governance and the Taming of Cyberspace*, Cambridge, Mass.: MIT Press. (p. 32)

NTIA, 2010. "Exploring the Digital Nation: Home Broadband Internet Adoption in the United States," <http://www.ntia.doc.gov/reports.html>.

Nuechterlein, J. E., and P. J. Weiser, 2005. *Digital Crossroads: American Telecommunications Policy in the Internet Age*, Cambridge, Mass.: MIT Press.

Odlyzko, A., 2010. "Bubbles, Gullibility, and Other Challenges for Economics, Psychology, Sociology, and Information Sciences." *First Monday* 15, no. 9. <http://www.uic.edu/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3142/2603>

Ou, G., 2008. "A Policy Maker's Guide to Network Management. The Information Technology and Innovation Foundation," <http://www.itif.org/index.php?id=205>.

Oxman, J., 1999. *The FCC and the Unregulation of the Internet*. Working paper 31, Federal Communications Commission, Office of Planning and Policy, Washington, D.C.

Partridge, C., 2008. "The Technical Development of Internet e-Mail," *IEEE Annals of the History of the Computing* 30, pp. 3-29.

Quarterman, J.S., 1989. *Matri Computer Networks and Conferences*, Bedford, Mass.: Digital Press.

Rosston, G., 2009. "The Rise and Fall of Third Party High Speed Access." *Information, Economics and Policy* 21, pp. 21-33.

Rosston, G., S. J. Savage, and D. Waldman, 2010. "Household Demand for Broadband Internet Service." SIEPR Paper 09-008. <http://siepr.stanford.edu/publicationsprofile/2109>.

Russell, A.L., 2006. Rough Consensus and Running Code and the Internet-OSI Standards War." *IEEE Annals of the History of Computing*, 28, pp. 48-61.

Saltzer, J.H., D.P. Reed, and D.D. Clark, 1984. "End-to-End Arguments in System Design," *ACM Transactions on Computer Systems* 2, pp. 277-288

.
An earlier version appeared in the *Second International Conference on Distributed Computing Systems* (April, 1981) pp. 509-512.

Savage, S. J., and D. Waldman, 2004, "United States Demand for Internet Access," *Review of Network Economics* 3, pp. 228-247.

Seamans, R., 2010. "Fighting City Hall: Entry Deterrence and New Technology Upgrades in Local Cable TV Markets." Working paper, NYU Stern School of Business, <http://pages.stern.nyu.edu/~rseamans/index.htm>.

Simcoe, T., 2007. "Delay and De jure Standardization: Exploring the Slow Down in Internet Standards Development," in (eds.) Shane Greenstein and Victor Stango, *Standards and Public Policy*, Cambridge, Mass.: Cambridge Press, pp. 260-297.

Simcoe, T., 2010. "Standard Setting Committees: Consensus Governance for Shared Technology Platforms," Working Paper, Boston University, <http://people.bu.edu/tsimcoe/index.html>.

Stranger, G., and S. Greenstein, 2007. "Pricing in the Shadow of Firm Turnover: ISPs in the 1990s." *International Journal of Industrial Organization*, 26, pp. 625-642.

Strover, S., 2001. "Rural Internet Connectivity." *Telecommunications Policy*, 25, pp. 331-347.

Von Burg, U., 2001. *The Triumph of Ethernet: Technological Communities and the Battle for the LAN Standard*, Palo Alto, CA: Stanford University Press.

Von Hippel, E., 2005. "Open Source Software Projects as User Innovation Networks," in (ed.) Joseph Feller, Brian Fitzgerald, Scott A. Hissam, Karim Lakhani, *Perspectives on Free and Open Software*, Cambridge, Mass.: MIT Press; pp. 267-278.

Von Schewick, B., 2010. *Internet Architecture and Innovation*, Cambridge, Mass.: MIT Press. (p. 33)

Waldrop, M., 2001. *The Dream Machine: J.C.R. Licklider and the Revolution that Made Computing Personal*, New York: Penguin.

Wallsten, S., 2009. "Understanding International Broadband Comparison," Technology Policy Institute, <http://www.techpolicyinstitute.org/publications/topic/2.html>

West, J., and S. Gallagher, 2006. Patterns of Innovation in Open Source Software, in (ed.) Hank Chesbrough, Wim Haverbeke, Joel West, *Open Innovation: Researching a New Paradigm*, Oxford University Press: Oxford, pp. 82–108.

Zittrain, J., 2009. "The Future of the Internet and How to Stop It, Creative Commons Attribution—Noncommercial Share Alike 3.0 license," www.jz.org.

Notes:

(1.) This is explained in considerable detail in Ou (2008).

(2.) There are many fine books about these developments, including Abbate (1999), Leiner et al. (2003), and Waldrop (2001).

(3.) Packet switching had been discussed among communications theorists since the early 1960s, as well as by some commercial firms who tried to implement simple versions in their frontier systems. As has been noted by others, the ideas behind packet switching had many fathers: Paul Baran, J. C. Likelider, Douglas Engelbart, and Len Kleinrock. There are many accounts of this. See, e.g., Quarterman (1989), Abbate (1999), Leiner et al. (2003), or Waldrop (2001).

(4.) These advances have been documented and analyzed by many writers, including, e.g., Quarterman (1989), Abbate (1999), Leiner et al. (2003), and Waldrop (2001).

(5.) An extensive explanation of TCP/IP can be found in many publications. Summarized simply, TCP determined a set of procedures for moving data across a network and what to do when problems arose. If there were errors or specific congestion issues, TCP contained procedures for retransmitting the data. While serving the same function of a postal envelope and address, IP also shaped the format of the message inside. It specified the address for the packet, its origin and destination, a few details about how the message format worked, and, in conjunction with routers, the likely path for the packet toward its destination. See, e.g., Leiner et al. (2003).

(6.) In part, this was due to a DOD requirement that all Unix systems do so, but it also arose, in part, because most Unix users in research environments found this feature valuable.

(7.) Moreover, a set of fail-safes had been put in place to make sure that one error message did not trigger another. That is, the system avoided the nightmare of one error

message generating yet another error message, which generated another and then another, thus flooding the system with never-ending messages.

(8.) The paper that defined this phrase is commonly cited as Saltzer et al. (1984).

(9.) As stated by Blumenthal and Clark (2001) in a retrospective look: “When a general-purpose system (for example, a network or an operating system) is built and specific applications are then built using this system (for example, email or the World Wide Web over the Internet), there is a question of how these specific applications and their required supporting services should be designed. The end-to-end arguments suggest that specific application-level functions usually cannot, and preferably should not, be built into the lower levels of the system—the core of the network.”

(10.) End-to-end also differed from most large computing systems at the time, such as mainframes, which put the essential operations in the central processing unit. When such computers became situated in a networked environment, there was little else for the terminals to do. They became known as *dumb* terminals. Contrast with end-to-end, which colloquially speaking, located *intelligence* at the edges of the system, namely, in the clients.

(11.) The economics behind this “surprise” is discussed in some detail in Greenstein (2011).

(12.) See, e.g., Aizcorbe et al. (2007), Doms and Forman (2005), or the discussion in Forman and Goldfarb (2006).

(13.) See the discussions in, e.g., Kahn (1995), Kahin and McConnel (1997), and Kesan and Shah (2001), or Greenstein (2010).

(14.) On many of the challenges during the transition, see Abbate (1999) or Kesan and Shah (2001).

(15.) For longer discussions about the origins and economic consequences, see, e.g., Mowery and Simcoe (2002a, b), and Greenstein (2011).

(16.) The term “mesh” is due to Besen et al. (2001).

(17.) See Friedan (2001).

(18.) For insights into the incentives to conduct traffic and come to peering agreements, see Besen et al. (2001) and Laffont et al. (2001, 2003).

(19.) For more on overlays, see Clark et al. (2006).

(20.) See the evidence in Rosston (2009).

(21.) This topic has received attention ever since the commercial Internet first began to blossom. See, e.g., Stover (2001) or Downes and Greenstein (2002). A summary can be found in Greenstein and Prince (2007). Also see the discussion in, e.g., the Federal Communications Commission (2010a).

(22.) See, e.g., Berners-Lee and Fischetti (1999), and Gilles and Cailliau (2000).

(23.) For more on this story, see Cusamano and Yoffie (2000).

(24.) Andrew Odlyzko maintains an accessible summary of studies and forecasts of data traffic at <http://www.dtc.umn.edu/mints/>.

(25.) Anderson and Wolff (2010) present a very accessible version of this argument.

(26.) Longer explanations for these events can be found in Greenstein (2007a, b).

(27.) For a longer discussion, see, e.g., Downes and Greenstein (2002, 2007).

(28.) For a discussion, see, e.g., Goldfarb (2004), Haigh (2007).

(29.) High usage could happen for a variety of reasons. For example, some technical users simply enjoyed being online for large lengths of time, surfing the growing Internet and Web. Some users began to operate businesses from their homes, remaining online throughout the entire workday. Some users simply forgot to log off, leaving their computers running and tying up the telephone line supporting the connection to the PC. And some users grew more experienced, and found a vast array of activities more attractive over time.

(30.) One survey of pricing contracts in May 1996 found that nearly 75 percent of the ISPs offering 28K service (the maximum dial-up speed at the time) offered a limited plan in addition to their unlimited plan. That dropped to nearly 50 percent by August. By March of 1997 it was 33 percent, 25 percent by January of 1998, and less than 15 percent by January of 1999. For a summary see Stranger and Greenstein (2007).

(31.) During the 1990s most cable companies sold access to the line directly to users but made arrangements with other firms, such as Roadrunner or @home, to handle traffic, routing, management, and other facets of the user experience. Some of these arrangements changed after 2001, either due to managerial preferences, as when @home lost its contract, or due to regulatory mandates to give users a choice over another ISP, as occurred after the AOL/Time Warner merger. See Rosston (2009).

(32.) Download speed may not reach the advertised maxima. In cable networks, for example, congestion issues were possible during peak hours. In DSL networks, the

quality of service could decline significantly for users far away from the central switch. The results are difficult to measure with precision.

(33.) The descriptive results were published in reports authored by staff at the NTIA. See NTIA (2010).

(34.) In addition to the surveys by Pew, also see, e.g., Savage and Waldman (2004), Rosston et al. (2010), and the summary of Greenstein and Prince (2007). For surveys of business adoption and its variance over geography, see Forman and Goldfarb (2006) and Forman et al. (2005).

(35.) See, e.g., Crandall (2005).

(36.) In many areas, fiber to the home was prohibitively expensive for almost all users except businesses, and even then, it was mostly used by businesses in dense urban areas, where the fiber was cheaper to lay. Fiber to the home has recently become cheaper and may become a viable option sometime in the future. See Crandall (2005).

(37.) OECD Broadband Portal, <http://www.oecd.org/sti/ict/broadband>, Table 1d.

(38.) OECD Broadband Portal, <http://www.oecd.org/sti/ict/broadband>, Table 1k. For a critique of the US standings in these rankings and insights into how to correct misunderstandings, see Wallsten (2009). Perhaps the biggest issue is the denominator, which is per capita. However, a subscriber tends to subscribe to one line per household. Since US average household size is much larger than average household size in other countries, the figure gives the false impression that US residences have less access to broadband than is actually accurate.

(39.) These come from Greenstein and McDevitt (2011).

(40.) For sources, see Greenstein and McDevitt (2009, 2011).

(41.) See e.g., Federal Communications Commission (2010b), or Marcus (2008).

(42.) See e.g., <http://www.wispa.org/>. Also see <http://wispassoc.com/>.

(43.) For interpretations of platform incentives, see e.g., Gawer (2009) or Evans, Haigu and Schmalensee (2006).

(44.) This is distinct from an engineering notion of a platform. The designers of the Internet deliberately built what they regarded as a computing platform. The inventors employed what they regarded as a balanced and standardized bundle of components to regularly deliver services. This balance reflected a long-standing and familiar principle in computer science. The inventors and DARPA administrators anticipated a benefit from

this design: others would build applications, though these inventors did not presume to know what those applications would do specifically.

(45.) See, e.g., Bresnahan and Greenstein (1999), or Dedrick and West (2001).

(46.) See, e.g., Bresnahan (1999).

(47.) See, e.g., Gawer and Cusumano (2002) or Gawer and Henderson (2007).

(48.) Sometimes Google retains many proprietary features, particularly in its search engine, which also supports a lucrative ad-placement business. Google takes action to prevent anyone from imitating it. For example, the caching, indexing, and underlying engineering tweaking activities remain hidden from public view.

(49.) See e.g., Cusumano and Yoffie (2000), Bresnahan and Yin (2007), or Bresnahan et al. (2011).

(50.) For an account of this decision, see Burgelman (2007).

(51.) For a book with this theme, see, e.g., Zittrain (2009), Lessig (1999), or Von Schewick (2010).

(52.) Anderson and Wolff (2010) or Zittrain (2009).

(53.) Once again, this case is explained in detail in Greenstein (2009b).

(54.) Numerous computer scientists and networking experts have pointed out that both Sprint and Cogent could have adjusted their routing tables in advance to prevent users from being cutoff. Hence, there is a sense in which both parties bear responsibility for imposing costs on their users.

(55.) It appears that Sprint's capitulation to its user base is, however, evidence that Sprint's management does not have the ability to ignore its users for very long.

(56.) This case is explained in detail in Greenstein (2009b).

(57.) For summary, see Adam Rothschild, December 2, 2010, <http://www.voxel.net/blog/2010/12/peering-disputes-comcast-level-3-and-you>.

(58.) This is especially true of the Internet Architecture Board and IETF, before it moved under the auspices of the Internet Society in 1992, where it remains today. See, e.g., Abbate (1999) or Russell (2006).

(59.) See Abbate (1999) for a history of the design of these protocols. See Partridge (2008) for a history of the processes that led to the development of email, for example.

(60.) Simcoe (2007, 2010) provides an overview of the operations at IETF and its changes as it grew.

(61.) For further discussion see Farrell and Simcoe, chapter 2 of this volume.

(62.) For example, see, e.g., Chiao et al. (2007).

(63.) These contrasts are further discussed in Greenstein (2009a).

(64.) See, e.g., Russell (2006).

(65.) The latter forbearance was deliberate. On the lack of interference in the design of the Ethernet, see von Burg (2001). On the design of 56K modems, see Augereau, Greenstein, and Rysman (2007). On the lack of regulation for network interconnection, see the full discussions in, e.g., Oxman (1999) or Kende (2000) or the account in Neuchterlein and Weiser (2005). More recent experience has departed from these trends, particularly in principles for regulating last-mile infrastructure. A summary of these departures is in Greenstein (2007b).

(66.) For a longer explanation of these origins, see, e.g., Kesan and Shah (2001) or Mueller (2002).

(67.) There is considerable writing about the growth of the production of Linux software and from a variety of perspectives. See, e.g., Dalle, David, Ghosh, and Wolak (2004), Lerner and Tirole (2002), VonHippel (2005), West and Gallagher (2006), or Arora and Farasat (2007). For an account and analysis of how many firms got on the Linux bandwagon, see, e.g., Dedrick and West (2001) or Fosfuri, Giarratana and Luzzi (2005). For further discussion see chapter 17 by Johnson, this volume.

(68.) For more on the history and operation of Apache, see, e.g., Mockus, Fielding, and Herbsleb (2005).

(69.) For more information, see Greenstein (2006).

(70.) For further discussion, see the chapter 3 by Hagiu, this volume.

(71.) The figures come from the website maintained by the Creative Commons, <http://creativecommons.org/>.

(72.) For one interesting account of the changing ratio of “suits to beards” at the IETF, see Simcoe (2007, 2010). For an account of the manipulation of hearings at the IEEE, see Mackie-Mason and Netz (2007).

(73.) Federal Communications Commission (2010a).

(74.) See Marcus (2008) and Marcus and Elixmann (2008).

(75.) The overinvestment in Internet infrastructure in the late 1990s had many causes. These are analyzed by, among others, Goldstein (2005), Greenstein (2007b), and Odlyzko (2010).

(76.) Statistics in the National Broadband Plan, FCC (2010a), seem to indicate that 3 percent of US homes subscribe to this service as of the end of 2009.

(77.) The economics behind the high cost of providing broadband in low-density locations is explained in detail in Strover (2001), Crandall (2005), and Greenstein and Prince (2007).

(78.) See Marcus (2008) and Marcus and Elixmann (2008).

(79.) For a summary, see Goldstein (2005), Nuechterlein and Weiser (2005), Greenstein (2010), and Federal Communications Commission (2010b).

(80.) See, e.g., Seamans (2010).

(81.) See, e.g., Baumol et al. (2006).

Shane Greenstein

Prof Shane Greenstein is Kellogg Chair in Information Technology at the Kellogg School of Management, Northwestern University, Evanston, USA

