Policy Analysis

No. 626

November 12, 2008

The Durable Internet
Preserving Network Neutrality without Regulation

by Timothy B. Lee

Executive Summary

An important reason for the Internet's remarkable growth over the last quarter century is the "end-to-end" principle that networks should confine themselves to transmitting generic packets without worrying about their contents. Not only has this made deployment of Internet infrastructure cheap and efficient, but it has created fertile ground for entrepreneurship. On a network that respects the end-to-end principle, prior approval from network owners is not needed to launch new applications, services, or content.

In recent years, self-styled "network neutrality" activists have pushed for legislation to prevent network owners from undermining the end-to-end principle. Although the concern is understandable, such legislation would be premature. Physical ownership of Internet infrastructure does not translate into a practical ability to control its use. Regulations are unnecessary because even in the absence of robust broadband competition, network owners are likely to find deviations from the end-to-end principle unprofitable.

New regulations inevitably come with unintended consequences. Indeed, today's network neutrality debate is strikingly similar to the debate that produced the first modern regulatory agency, the Interstate Commerce Commission. Unfortunately, rather than protecting consumers from the railroads, the ICC protected the railroads from competition by erecting new barriers to entry in the surface transportation marketplace. Other 20th-century regulatory agencies also limited competition in the industries they regulated. Like these older regulatory regimes, network neutrality regulations are likely not to achieve their intended aims. Given the need for more competition in the broadband marketplace, policymakers should be especially wary of enacting regulations that could become a barrier to entry for new broadband firms.

Timothy B. Lee, an adjunct scholar at the Cato Institute, is pursuing a Ph.D. in computer science at Princeton University.

CATO INSTITUTE

Too much centralization and bureaucracy is detrimental to innovation.

Introduction

The 1970s saw two revolutions that would transform the American economy for decades to come. One was the early development of the Internet. The other was a wave of deregulation that freed the nation's transportation and communications infrastructure from micromanagement by federal bureaucracies.

Each of those revolutions was tied to an intellectual tradition that has profoundly shaped the modern world. In the 1980s, the Internet was one network among many, and most of its competitors were built on proprietary standards. Partisans for the Internet tended to be partisans for open technologies more generally. As the Internet has emerged as the undisputed winner of the networking wars, it also became the poster child for "openness," the now-dominant ideology of Silicon Valley.

Similarly, the deregulations of the 1970s were brought about by a sea change in scholarly attitudes toward government regulation. Public policy scholars in the early 20th century had imagined that neutral bureaucrats could manage the economy and society. That naïve optimism gave way to a more sophisticated and skeptical view of the regulatory process in the decades after World War II. Economists began to suggest that regulatory processes were vulnerable to "government failures" akin to the market failures often cited to justify government regulations. Scholars articulated theories of "regulatory capture" in which regulated industries manipulated the regulatory process for their own benefit. And they began to recognize the frequency with which regulatory schemes produce harmful, unintended consequences.

In "The Broadband Debate: a User's Guide," Columbia law professor Tim Wu dubbed these two schools of thought the "openists" and the "deregulationists," respectively.[1] The network neutrality debate has put the heirs of these traditions on a collision course. Each camp views the other as a threat to the gains of the last quarter century. Openists worry that the remnants of the

Bell system will regain control over the nation's communications infrastructure and transform the Internet into a proprietary network. Deregulationists, on the other hand, worry that Washington bureaucrats will gain control over the Internet, returning the country to the bad old days when government bureaucrats, not market forces, determined the shape of communications markets.

These two movements have come to regard themselves as implacable foes, but they have more in common than they like to admit: they share the fundamental insight that too much centralization and bureaucracy is detrimental to innovation. But each is convinced that the other's agenda will bring about these unfortunate circumstances. Each camp has sometimes overstated its case and failed to take the other side's concerns seriously. And each camp has a great deal to learn from the other.

The openist camp includes Internet pioneers like World Wide Web inventor Timothy Berners-Lee,[2] who is intimately familiar with the prerequisites for online innovation. It would be a mistake to dismiss too lightly this camp's concerns about the problems that could be created by network discrimination. The deregulationists include prominent economists such as Alfred Kahn, who oversaw the deregulation of the airline industry under President Jimmy Carter. Kahn possesses a deep understanding of the unintended consequences of government regulation. Ignoring his concerns about the unintended consequences of government regulation would be equally misguided.[3]

"Network neutrality" has been given many meanings, but the core dispute is over whether network owners will alter the Internet's end-to-end architecture. Openists fear this outcome while some deregulationists welcome it. Other deregulationists flatly deny that the end-to-end principle has ever been the norm on the Internet. But in fact, the end-to-end principle has been the central organizing principle of the Internet for a quarter century. And both sides overestimate the power of the network owners. The natural inertia of the Internet's architecture, together with the vigi-

lance and technical skill of the online community, is likely to provide an adequate counterweight to network owners' efforts to transform the Internet into a proprietary network, regardless of the actions of government regulators. Network owners who try to profit from discriminatory practices will encounter stiff resistance from an army of tech-savvy users who rapidly develop and disseminate countermeasures and workarounds. Network owners will find that they lack the leverage to effectively control the behavior of online firms and users and that efforts to limit the activities of their own customers are financial and public-relations disasters. Network owners who try to construct a "walled garden" of proprietary applications and content are likely to be similarly disappointed, as proprietary services fail to keep pace with the open Internet. ISPs are likely to respect network neutrality not because they want to but because economic and technological constraints leave them little choice.

Concerns that network owners will undermine free speech online are particularly misguided. Network owners have neither the technology nor the manpower to effectively filter online content based on the viewpoints being expressed, nor do profit-making businesses have any real incentive to do so. Should a network owner be foolish enough to attempt large-scale censorship of its customers, it would not only fail to suppress the disfavored speech, but the network would actually increase the visibility of the content as the effort at censorship attracted additional coverage of the material being censored.

The openists have a tendency to underestimate the unintended consequences that can occur when governments regulate. History suggests that regulatory efforts to protect the customers of major infrastructure owners often end badly. The first such effort was the creation of the Interstate Commerce Commission in 1887. The language of the Interstate Commerce Act was strikingly similar to the network neutrality language being considered today. The ICC's backers touted it as a way of protecting the public from abuses by the railroads, but in practice it reduced competition in the railroad industry, effecting transfers of wealth from the general public to the railroads and other politically connected groups. Over the course of the 20th century, the ICC transformed virtually the entire surface transportation industry into a government-run cartel. This and other examples suggest that policymakers should be extremely cautious about enacting new regulations when nonregulatory approaches might achieve the same goals.

If there's one thing that almost all sides of the network neutrality debate agree on, it's that there is inadequate competition in the broadband marketplace. Given that consensus, openists should think twice about demanding new regulatory regimes that could create barriers to entry for new market entrants. Complying with regulatory regimes requires the services of lawyers, lobbyists, accountants, and other highly paid professionals. Every dollar spent on these activities is a dollar that cannot be spent on R&D or new infrastructure. Regulations designed with today's technologies in mind could in practice bar new entrants with innovative business models and technologies. Congress should therefore be reluctant to impose regulations on a sector of the economy that has, until now, been largely free to do business without government regulation.[4]

The existence of significant constraints on the power of individual network owners and the risks of unintended consequences suggest that enacting prospective network neutrality regulation would be premature and probably counterproductive. There is little danger that network owners will fundamentally transform the Internet's architecture, and so it would be unwise for policymakers to enact new regulations to deal with vague or speculative threats.

The Internet and End-to-End

The Internet owes its extraordinary success to a set of technical principles that have been implicit in its design since it began life as an

ISPs are likely to respect network neutrality not because they want to but because economic and technological constraints leave them little choice.

experimental network called the ARPANET in 1969. The most important of these principles is the "end-to-end" principle, which holds that computer networks should be decentralized, with most of the "intelligence" of the network residing on computers at the network's endpoints, rather than with routers at the core of the network. The end-to-end principle gave the Internet important technical and economic advantages that helped it to outpace its rivals and become the world's dominant communications network.

The military's Advanced Research Project Agency (ARPA) was tasked with funding scientific research that might have military applications. During the 1960s, it provided significant funding to computer-science researchers. One of the most important projects it funded was an experimental packet-switched network called the ARPANET. The networks already in existence at the time of the ARPANET's founding were centrally managed, special-purpose networks. The telephone system, for example, was optimized for carrying phone calls and little else. Indeed, AT&T strongly discouraged users from using it for other purposes. Adding a new capability (or "functionality") to the phone network was a costly proposition because it often required a comprehensive overhaul.

In contrast, ARPANET was designed to accommodate a variety of different applications without modifications to the network. By transmitting generic packets, rather than phone calls or telegraph messages, it left the sender and receiver free to decide for themselves what kind of information would be transmitted and how that information would be interpreted by the recipient. That gave the system unprecedented flexibility. The network's designers initially envisioned the creation of file-transfer and remote-login applications (the predecessors of today's FTP and Telnet protocols[5]), but they wanted to make it easy for users to develop additional applications that were not envisioned by the network's creators.

That foresight paid off in 1972, when Ray Tomlinson, an engineer at the firm that built

and managed the ARPANET, developed the first e-mail program. E-mail rapidly became the Internet's "killer app." Prior to the introduction of e-mail, ARPANET was little more than an interesting research project. But once the head of ARPA, Stephen Lukasik, started using e-mail to manage day-to-day ARPA business, it became an indispensable communications tool for military researchers.[6] Indeed, a study commissioned just a year after the debut of e-mail found that it already accounted for three-quarters of all traffic on the ARPANET.[7]

Crucially, Tomlinson didn't need to ask the permission of the ARPANET's operators to develop his simple e-mail program, nor did ARPANET users require anyone's approval to begin using it. All that was required to add e-mail functionality to ARPANET was for the sender and receiver to each have compatible software installed on their computers.

ARPANET differed from the modern-day Internet in at least one important respect: the network was responsible for guaranteeing that every packet made it from source to destination. In the event of congestion or other problems, the network would store the packet and resend it at a later time, ensuring that it would eventually be delivered. That design worked well on a homogenous network with extremely limited computing power at the endpoints. ARPANET's reliability guarantees reduced the burden on the programmers developing new applications because they didn't have to worry about lost packets.

However, as packet-switched networks became more popular, it became clear that the ARPANET design had a serious weakness: the ARPANET protocols were not well suited for combining heterogeneous networks. Because ARPANET applications depended on the network's reliability guarantees, all parts of the network had to be designed to support them. Researchers experimenting with new types of networks (such as wireless) in which packet loss was more common became convinced that a new set of protocols was needed to join together networks with diverse operating characteristics.

ARPANET was designed to accommodate a variety of different applications without modifications to the network.

4

TCP/IP

The result, developed over the course of the 1970s, was the TCP/IP protocol suite (that is, the collection of protocols) that forms the foundation of today's Internet.[8] TCP/IP achieved unprecedented flexibility by shifting even more responsibility to the network's endpoints. Unlike the ARPANET, TCP/IP-based networks offer no guarantees about reliable packet delivery. When an IP-based router encounters congestion or other problems, it simply drops any packets it can't deliver in a timely fashion, making it the responsibility of the sender to notice that a packet hasn't been acknowledged and re-send it.[9]

A seminal 1984 paper by three MIT computer scientists made explicit the end-to-end design principles implicit in TCP/IP. In their paper, "End-to-End Arguments in System Design," J. H. Salzer, D. P. Reed, and D. D. Clark pointed out that placing functionality at the endpoints of a network, rather than within the network itself, could both lower costs and increase the flexibility of the system.[10] The developer of any given application will have a better idea of the appropriate level of functionality than will the designers of a network that might be used by many different applications, so the end-to-end approach avoids adding unnecessary (and costly) functionality to a network that won't be used by many applications. Although the original end-to-end paper didn't focus on TCP/IP networks specifically, the Internet has since become the paragon of the principles described in the paper. On a TCP/IP network, a router's only responsibility is delivering individual packets to their destination. Decisions about what to do with those packets are made by applications running at the network's endpoints.

The advantages of this decentralized approach to network design is best illustrated by the birth of the World Wide Web, an application that many now regard as synonymous with the Internet itself but which was not invented until 1990. The Web was developed by computer scientist Timothy Berners-Lee at the European physics laboratory CERN. Berners-Lee received limited support from his superiors,

and as a result the initial effort to develop the World Wide Web was a shoestring affair with only a handful of collaborators.[11] Luckily, the Internet's decentralized design meant that no special modifications to the Internet's architecture—or even permission from its operators— were needed to deploy the Web around the world. As soon as Berners-Lee completed work on the first Web server, anyone who had an Internet connection and a copy of his Web browsing software could access it. Thanks to the end-to-end principle, Berners-Lee and dozens of others were able to launch new Internet applications that could reach a worldwide audience at minimal cost.

The Limits of Closed Networks

The current market for software on mobile phones provides a stark contrast to this happy state of affairs. In many cases, mobile applications can only be brought to market with the explicit permission of the major wireless operators. Tim Wu has argued that developing software for mobile phones can be extremely frustrating, because on most mobile platforms, developers have to spend as much time dealing with the carriers' bureaucratic approval process as they do actually developing their software. Wireless carriers routinely impose elaborate testing requirements, demand a cut of application developers' revenues, and even ban software functionality that might conflict with their existing business models. The result is an anemic market for phone-based software.[12]

Wu overstates his case in some respects. It is possible to get smart phones that are relatively free of carrier restrictions, although those phones tend to be unsubsidized by the carriers and therefore are significantly more expensive. Phones based on Microsoft's PocketPC platform, for example, place relatively few restrictions on software development. And recent developments including the introduction of the iPhone,[13] the unveiling of Google's Android mobile operating system,[14] and Verizon's announcement that it will make its network more open,[15] suggest that competitive pressures may continue to push wireless carriers toward greater openness. As Wu himself acknowledges,

Placing functionality at the endpoints of a network, rather than within the network itself, both lowers costs and increase the flexibility of the system.

the predominance of closed cellular platforms is not necessarily an argument for government regulation.

But Wu's paper powerfully illustrates the downsides of closed networks. At present, it is not as easy as it should be for a smart programmer to release innovative new mobile software on a shoestring budget the way Berners-Lee did with the Web two decades ago. Developers that wish to release software for the most widely deployed cell phone platforms must run the gauntlet of the major carriers' approval processes. There are probably talented software developers who are being excluded (and consumers being denied innovations) because of these barriers. We should, therefore, be concerned at the prospect of the Internet's becoming more like a closed wireless network.

Recently, legal scholars have begun using the term "network neutrality" to denote a principle roughly equivalent to the end-to-end principle. The concept of "network neutrality" is rarely defined precisely, and network neutrality advocates sometimes disagree among themselves about precisely which activities violate it. For purposes of clarity, this paper will mostly use the slightly narrower concept of "end-to-end." There are a few examples of network policies that arguably violate network neutrality but not the end-to-end principle. These are considered explicitly later in this paper. But the end-to-end principle is a precise technical concept that encompasses the great majority of the behaviors that concern network neutrality advocates.

The Internet succeeded because, thanks to the end-to-end principle, it harnessed the power of decentralized innovation, allowing anyone to add functionality to the network without centralized decisionmaking. Today, the importance of the end-to-end principle has become the conventional wisdom, but there are dissenters from this networking orthodoxy. In recent years, several important economists, legal scholars, and technologists have argued that the emergence of new applications and growing Internet congestion have strengthened the case for "smarter" networks that give differential treatment to packets based on their contents.

> **The Internet succeeded because, thanks to the end-to-end principle, it harnessed the power of decentralized innovation.**

The Underwhelming Case for Network Discrimination

A representative example of the argument for network discrimination was put forward by economists Robert Hahn and Scott Wallsten of the (now-defunct) AEI-Brookings Joint Center for Regulatory Studies.[16] Hahn and Wallsten view network neutrality regulation as a kind of price regulation, requiring that "last mile" broadband providers charge content providers a price of zero for access to their networks. They suggest that under a network neutrality regime, application developers would have no incentive to "take into account potential congestion costs of bandwidth-intensive applications."

This argument is based on a flawed understanding of the Internet's basic structure. The Internet is a collection of thousands of networks that agree to carry one another's traffic. Any pair of networks that wish to exchange traffic negotiate an agreement specifying the terms of interconnection. If one network is significantly larger than the other, the smaller network will typically pay the larger network for connectivity, an arrangement known as "transit." If two networks are roughly equal, they will typically carry each others' traffic without charge, an arrangement known as "settlement-free peering."[17] Because these agreements are negotiated in the context of a competitive market, they tend to reflect the full cost to each network of carrying the other's traffic.

The price that an Internet firm like Google pays for bandwidth includes the costs of securing "upstream" connectivity to other networks. The costs of delivering traffic to a "last mile" broadband provider like Comcast or Verizon is implicitly included in the price Google pays for connectivity. Hence, Hahn and Wallsten's suggestion that network neutrality allows Internet companies to "use [network owners'] property for free" is mistaken. Network owners do not receive direct payments from all of the parties

whose data they carry, but the network of consensual interconnection agreements that binds the Internet together ensures that each Internet user pays a fair share of the total costs of running the network.

With thousands of network owners and hundreds of millions of users, it would be prohibitively expensive for every network to charge every user (or even every online business) for the bandwidth it uses. Transaction costs would absorb any efficiency gains from such an arrangement. It would make no more sense than an automobile manufacturer requiring its customers to make separate payments to the manufacturers of every component of a new automobile. One of the services an ISP provides to its customers is "one stop shopping" for Internet connectivity. This arrangement has important economic advantages and is unlikely to change in the foreseeable future.

Application and Content Discrimination

Another prominent critic of network neutrality is Christopher Yoo, a law professor at Vanderbilt University. Yoo correctly argues that we should be cautious about enacting legislation that might foreclose beneficial evolution of networking technologies. But Yoo goes beyond that general note of caution to offer some specific arguments for abandoning the end-to-end principle. Yoo's arguments on this score are unpersuasive.

Yoo contends that discriminating among bandwidth-hogging applications could be an effective way to minimize congestion.[18] He suggests that by charging users different prices depending on the types of applications they wish to use, or prohibiting the use of certain bandwidth-hogging applications altogether, ISPs could avoid network gridlock and improve the experience of ordinary users who use standard Internet applications like the Web and e-mail.

There are two major problems with this approach to managing congestion. One is that novice users are likely to find it confusing. There are thousands of Internet applications, including online games, business applications,

and social networking websites. The average Internet user has no clear sense of the type or amount of data any given application generates. Any pricing policy complex enough to distinguish among the many categories of Internet content is likely to be incomprehensible to most customers.

Tech-savvy users would present an even bigger headache for a network owner with a discriminatory pricing strategy. If different types of data were billed at different rates, users would have a powerful financial incentive to camouflage their high-priced bits to look like lower-priced bits. That would spark a technological arms race in which the ISP developed more sophisticated filtering technology and users developed better evasion techniques. Network owners would almost certainly lose this arms race, but not before spending millions of dollars on unnecessary hardware and software.

If additional measures are needed to control congestion, it's likely to be far more effective to instead impose content-neutral restrictions on bandwidth consumption. These could take the form of bandwidth caps, metering, or limits on average throughput. Any of these strategies could relieve congestion without the problems of discriminatory traffic filtering.[19]

Quality of Service

Some applications, such as Internet telephony and online gaming, are extremely sensitive to delays in packet delivery (known as latency). Yoo suggests that it could be beneficial for networks to give packets from these applications higher priority than packets from applications like the Web and e-mail that are less sensitive to latency.[20] Most Internet applications, including the Web and e-mail, are not significantly disturbed by short delays in packet delivery. However, some applications can be significantly degraded by what network engineers call "jitter," or random delays in packet delivery. Probably the most significant latency-sensitive application is Internet telephony. An occasional one-second delay in packet delivery can dramatically degrade the quality of phone

The Internet is a collection of thousands of networks that agree to carry one another's traffic.

7

The Internet was deliberately designed without a centralized authority that could allocate bandwidth to favored applications.

calls using a VoIP application (Voice over Internet Protocol) such as Vonage or Skype.

"Quality of service" (QoS) technologies attempt to guarantee a latency-sensitive application such as Vonage a minimum amount of bandwidth at all times. This obviously involves routing packets from certain (latency-sensitive) applications in preference to packets from other applications, and as such it would appear to be a violation of the end-to-end principle.

Yoo contends that the need for QoS guarantees is a strong argument for relaxing the end-to-end principle. But Ed Felten, a computer scientist at Princeton University, offers a couple of reasons why QoS guarantees may not be as necessary as they seem at first glance. A latency-sensitive application can sometimes be converted into a non-latency-sensitive application through clever engineering. For example, streaming video is latency-sensitive, but newer video applications such as YouTube employ buffering, so that on a fast enough network connection they almost always display the video smoothly. Second, QoS guarantees are not needed on a network with a lot of spare capacity. If an application's bandwidth needs are significantly less than the average bandwidth available on the network, short-term fluctuations in available bandwidth may not cause problems because the throughput may never drop below the application's minimum rate. Felten suggests that some fast networks may have reached this point for voice applications.[21]

Quality-of-service guaranties may prove so expensive to implement that it would be more cost-effective to focus on increasing total capacity instead. QoS guarantees are hard to implement on a heterogeneous network like the Internet. From 1998 to 2001, a group of researchers associated with the Internet2 project conducted a series of experiments with QBone, an experimental QoS architecture. In 2002, they released a report concluding that QBone suffered from "poor incremental deployment properties, intimidating new complexity for network operators, missing functionality on routers, and serious economic challenges." They argued

that the costs of QoS architectures are higher than the benefits and would "threaten the scalability and flexibility of the Internet."[22]

The fundamental problem is economic as much as it is technical: introducing QoS features makes network interconnection much more complicated. It's relatively easy to implement QoS guarantees on an integrated network owned entirely by one network provider, because the network can have centralized management infrastructure that allocates the necessary bandwidth to each application. But Internet traffic almost always traverses more than one network, and a QoS guarantee for half of a network path isn't worth much. The Internet was deliberately designed without a centralized authority that could allocate bandwidth to favored applications. Nor is there anything resembling a billing infrastructure that would allow applications to purchase guaranteed bandwidth on other networks.

If effective QoS technologies are developed, they are likely to be implemented in a decentralized manner that is consistent with the spirit of the end-to-end principle. For example, one of the most prominent schemes for packet prioritization is DiffServ, developed in the late 1990s.[23] Under this scheme, network endpoints mark each packet with one of a small number of priority classes. When routers encounter congestion, they drop lower-priority packets before higher-priority ones. To prevent cheating, routers at network boundaries limit the number of high-priority packets the network will accept per user, reclassifying or dropping packets to enforce pre-existing limits on the number of high-priority packets the network will accept.[24]

A more ambitious proposal was laid out in a recent paper by Lawrence G. Roberts, the man who led the original ARPANET project, and two other researchers. It proposes a QoS architecture for the Internet that would allow any network endpoint to request bandwidth guarantees using a standardized protocol, with intermediate routers indicating whether they have the spare capacity to guarantee the requested bandwidth.[25]

These designs may be inconsistent with the

end-to-end principle narrowly conceived, but they are consistent with the end-to-end principle in the sense that they leave network endpoints with the ultimate authority to decide which packets should get priority treatment. Under these schemes, networks do *not* attempt to classify packets based on their contents or prioritize based on the network owner's judgments about which applications or content merit priority treatment. Rather they allow users and application developers to decide which applications are latency-sensitive.

There would be no real advantage, and considerable disadvantages, to having network owners try to recognize latency-sensitive traffic based on packet contents. First, such a scheme would undermine one of the Internet's core strengths: the ability of new applications to be deployed without consulting the hundreds of companies that manage various parts of the Internet. If network owners adopted lists of latency-sensitive applications that would receive higher priority, a company launching a new, latency-sensitive application would need to lobby dozens of network operators for inclusion on their lists. Second, if packets were prioritized based on the type of application, applications that didn't make the cut would have a strong temptation to boost performance by camouflaging their traffic so that it looked like the traffic of a high-priority application, once again sparking an unnecessary arms race. If prioritization is the goal, it makes more sense to allow users themselves (specifically, the applications they choose to install and run) to explicitly mark the priority of their packets rather than having the network try to guess the appropriate values.

Editorial Filters

Yoo argues that the explosion of content on the Internet has made it necessary for "telecommunications networks to exercise editorial control."[26] Analogizing the Internet to a cable television network, he suggests that network owners need the ability to decide which websites their customers visit for the same reasons that cable operators decide which channels to carry. He suggests that such content selection by network owners is analogous to the editorial policies of websites such as Google or *Sports Illustrated*. Yoo worries that consumers will be harmed if they cannot be provided with "editorial filters."

This critique seems to miss the fundamental difference between traditional cable networks and the Internet. Analog cable networks broadcast all of their channels over the wire simultaneously. As a consequence, there is a limit to the number of cable channels that can be made available to the user. It is therefore unavoidable that someone will decide which cable channels will be provided. In contrast, Internet content is transmitted only upon user requests. As a result, there's no need for ISPs to pick and choose among Internet content. They can make all the applications and content on the Internet available, and let the user choose.

Yoo is right about the importance of editors to filter the avalanche of information available on the Web. But he misunderstands the fundamental division of labor between the routers in the core of the network and servers at the endpoints. No large network owner could build a filtering regime that would satisfy each of its millions of customers. At the same time, if network owners respect the end-to-end principle, users can choose from among the thousands of filters already available on the Internet. Websites like Google, Digg, *Sports Illustrated*, or ICanHasCheezburger help users find content they're interested in and weed out the rest. There are no good reasons for ISPs to try to displace this abundance of filtering options, and good reasons to hope they don't.

The end-to-end principle ensures that end users have maximum control over their Internet experience. Deviations from end-to-end will generally reduce user autonomy by substituting the network owner's judgment for the user's own judgment. The arguments in favor of doing this are unpersuasive. But Yoo and other scholars have also argued that the end-to-end principle has *already* been abandoned online. We turn to these arguments next.

There would be no real advantage, and considerable disadvantages, to having network owners try to recognize latency-sensitive traffic based on packet contents.

9

Is the Internet Neutral Now?

Yoo cites the "emergence of beneficial practices, such as backbone [i.e., settlement-free] peering, content delivery networks like Akamai, network-based spam filtering, and blocking websites known to be sources of viruses"[27] as examples of current networking practices that violate the end-to-end principle. But with one possible exception, these examples do not offer compelling arguments for relaxing the end-to-end principle. To understand why, it is important to recognize that the end-to-end principle constrains only the routers "inside" the network that are responsible for routing the packets of other computers. The end-to-end principle does *not* constrain the behavior of network endpoints, which never handle any packets other than their own.

As discussed above, settlement-free peering is an arrangement in which two networks agree to carry each others' traffic without charge. This typically occurs when the networks are of roughly equal size, and so the benefits of peering to each side are approximately equal. On the other hand, when networks of unequal size connect, the smaller network will often be required to pay the larger network to carry its traffic.

Yoo's concern seems to be that the different financial treatment of large and small networks violates the principle that all traffic be treated equally. But it must be remembered that the end-to-end principle, and network neutrality more broadly, are focused on the technical, rather than contractual, behavior of network owners. The end-to-end principle requires that a network's routers give equal treatment to all packets that traverse the network. It has nothing to say about the prices networks charge each other for interconnection. Backbone peering simply doesn't violate network neutrality or the end-to-end principle, if those terms are properly understood.

Akamai

A content delivery network consists of thousands of servers distributed around the world that cache frequently-accessed content on behalf of clients. For example, CNN might arrange for Akamai to host its video content. Instead of downloading videos directly from CNN's web server, the user's web browser downloads the content from an Akamai-owned server close to the user's location.

Yoo's argument regarding content-delivery networks stems from a misunderstanding about the nature of those networks. The word "network" has a number of distinct meanings in computer science, and a content-delivery network is not a "network" in the same sense that the Internet is a network. It is a network only in the more general sense of a group of computers working together to achieve a common purpose. Given the confusing terminology, it's understandable that Yoo would assume that "intelligence in the core of the network" is required for Akamai to work properly. But in practice, Akamai's servers communicate via ordinary TCP/IP connections, and Internet routers route Akamai packets exactly the same way they route any other packets.

To understand how Akamai manages this feat, it's helpful to know a bit more about what happens under the hood when a user loads a document from the Web. The Web browser must first translate the domain name (e.g., "cato.org") into a corresponding IP address (72.32.118.3). It does this by querying a special computer called a domain name system (DNS) server. Only after the DNS server replies with the right IP address can the Web browser submit a request for the document. The process for accessing content via Akamai is the same except for one small difference: Akamai has special DNS servers that return the IP addresses of different Akamai Web servers depending on the user's location and the load on nearby servers. The "intelligence" of Akamai's network resides in these DNS servers.

Because this is done automatically, it may seem to users like "the network" is engaging in intelligent traffic management. But from a network router's perspective, a DNS server is just another endpoint. No special modifications are needed to the routers at the core of

> **Backbone peering simply doesn't violate network neutrality or the end-to-end principle.**

the Internet to get Akamai to work, and Akamai's design is certainly consistent with the end-to-end principle.

Spam and Viruses

The same point applies to Yoo's example of spam filtering. To see why once again requires a brief discussion of Internet architecture. From an architectural perspective, e-mail servers are network endpoints, just like Akamai's Web and DNS servers. Internet routers route e-mail packets the same way they route any other type of packet. Although e-mail service is sometimes bundled together with Internet access, there is no necessary connection between the two. Indeed, many Internet users use third-party e-mail access that is not affiliated with their ISP.

One advantage of this arrangement is that users of neutral networks can choose third-party e-mail providers if they are dissatisfied with the e-mail service provided by their own ISP. In contrast, filtering e-mail at the network level imposes one anti-spam policy on every user, whether or not they appreciate this "service." Some ISPs *do* engage in network-level spam filtering, but this activity is not essential to anti-spam efforts and is arguably counter-productive.[28]

The end-to-end principle does not preclude an ISP from offering spam filtering services on its own mail server. But it does require network owners not to interfere with users who wish to use a mail server provided by a third party. This ensures that users who are dissatisfied with the anti-spam policies of their ISP's own e-mail service can choose another one.

Virus-infested websites are a rare case where a strong argument can be made for deviating from the end-to-end principle. But it's not difficult to draw a principled distinction between efforts to combat viruses and most other deviations from end-to-end. Anti-virus efforts are typically designed to protect users against malicious strangers. That seems fundamentally different from run-of-the-mill violations of network neutrality that prioritize some legitimate users or applications over others. In any event, the most prominent network neutrality

legislation, sponsored by Sens. Olympia Snowe (R-ME) and Byron Dorgan (D-ND), would have prohibited only interference with "lawful content, application or service."[29] Viruses would likely be considered illegal applications under this definition. The need to combat the spread of viruses, therefore, does not seem to be a compelling argument against leading network neutrality proposals.

Misreading RFCs

Economists Robert Hahn and Robert Litan have also claimed that adherence to the end-to-end principle is far from universal. In a paper for the AEI-Brookings Joint Center on Regulation they argued that the Internet does not follow the end-to-end principle and never did.[30] Their major evidence is found in technical documents called "requests for comments" (RFCs) that define basic Internet protocols. Hahn and Litan argue that at least four RFCs appear to countenance nonneutral routing of Internet packets. However, closer inspection of these documents gives a very different picture. One document,[31] written in 1994, does not describe the existing TCP/IP protocols, but a "proposed extension to the Internet architecture" that has not been widely adopted for public Internet connectivity.[32] A second, written in 1974,[33] advises host machines implementing the TCP protocol to "treat incoming packets with higher priority than outgoing packets." But remember that the end-to-end principle constrains the behavior of routers in the core of a network, *not* hosts at its end points. The third, published in 1981,[34] specifies that packets in the IP protocol should include a field for priority and that this field could be used for prioritizing packets. However, the RFC doesn't specify how routers should use this information, and the field is generally ignored by modern Internet routers. Finally, Hahn and Litan cite another 1981 paper by Internet pioneer Vinton Cerf that did indeed describe a nonneutral networking scheme.[35] However, the document concerns AUTODIN, an early *alternative* to the TCP/IP protocol suite that never caught on. That a failed early competitor to TCP/IP did not observe the end-to-

Large-scale violations of the end-to-end principle have certainly been rare and have almost always generated controversy.

end principle is certainly not evidence that the modern Internet violates it.[36]

It would be overstating the case to claim that the end-to-end principle has never been violated. But large-scale violations of the end-to-end principle have certainly been rare and have almost always generated controversy. Neutral treatment of packets by "dumb" networks has been the norm for a quarter century, and there are good reasons to preserve that arrangement. Respecting the end-to-end principle ensures that end-users are in control of their Internet experience, and it provides a fertile environment for online innovators, who are able to quickly and easily reach a global audience with new content and applications.

There is a widespread assumption on both sides of the network neutrality debate that the Internet's end-to-end architecture is quite fragile. Many people believe that network owners have broad powers to reshape the Internet, or at least their own customers' experience of it. But this is far from true.

Customers Gone Wild: Why Ownership Doesn't Mean Control

Debates over regulatory policy are replete with claims that network owners will—a few say "should"—allow, prohibit, promote, or discourage a variety of applications, devices, and content on their networks. They can, it is imagined "speed up" favored applications and "slow down" disfavored applications, make some content more prominent than others on users' screens, and tilt the direction of the online conversation in ways that are congenial to network owners. Proponents of network neutrality regulations warn that this outcome will lead to a less innovative, less useful, and less democratic Internet. Some opponents of regulation welcome it, suggesting that deviations from end-to-end can increase the efficiency of the network, reduce congestion, and accomplish other worthwhile goals. But hardly anyone questions whether companies would be able to undermine the Internet's end-to-end architecture.[37]

Yet example after example suggests that, in practice, the owners of large, open technological platforms have only limited control over the use of those platforms. As Apple has discovered with its ongoing attempts to lock down its iPhone platform,[38] customers cannot be counted on to passively accept artificial limitations imposed by platform owners. To the contrary, customers actively resist such restrictions, and in many cases, platform owners find themselves almost powerless to prevent it. An example will help to illustrate this point.

09 F9 11 02 9D 74 E3 5B D8 41 56 C5 63 56 88 C0

In early 2007, hackers extracted a previously secret "processing key" that could be used to unscramble commercial HD-DVD and Blu-Ray discs without the permission of copyright owners. This 128-bit key, which can be represented in hexadecimal notation as "09 F9 11 02 9D 74 E3 5B D8 41 56 C5 63 56 88 C0," began to circulate on the Internet. In April, the "licensing authority" in charge of the copy protection built into Blu-Ray and HD-DVD discs began sending threatening letters to website operators demanding that the key be removed. This effort backfired, as the threatening letters generated more publicity for the key, and more people began hosting the key on their websites.[39]

Soon the key appeared on Digg, a popular online news site. Digg is unlike most other news sites because its content is created and selected almost entirely by Digg readers. They submit stories to the site and vote on which stories should appear on the front page. The results reflect the quirky tastes of the site's largely young, male, tech-savvy audience, with a mixture of technology and gadget news, politics, celebrity gossip, and sophomoric humor. Like other sites, Digg received legal demands that the keys be removed from the site. In a blog post on May 1, Digg CEO Jay Adelson announced that Digg would comply with the demands.[40]

Digg users regarded this decision as censorship. Over the next eight hours, thousands of

Digg users began submitting stories containing the key and voting those stories onto the site's front page. At one point, Digg's entire home page was filled with stories about the key. Digg's staff attempted to delete the stories but they were simply unable to keep up. Nor would software filters have been an effective solution because users began posting the key in a variety of formats that couldn't easily have been blocked by filters. One user picked up his guitar and turned the key into a song. Another user registered the key as a domain name and posted a link to that. Others posted images and videos featuring the key. T-shirts with the key printed on them were offered for sale.

Just eight hours after Adelson's initial announcement, Digg threw in the towel. In a later blog post that same day, Digg founder Kevin Rose wrote that, "after seeing hundreds of stories and reading thousands of comments, you've made it clear. You'd rather see Digg go down fighting than bow down to a bigger company."[41] Rose put a brave face on the situation, but the reality is that Digg had no real choice. Its users had demonstrated their determination to keep the key on Digg's front page. Digg's users vastly outnumbered its employees, and their ingenuity and determination were no match for any software Digg could reasonably have developed.

Of course, Digg could have overhauled the site to make it more like a more traditional news site, in which Digg employees reviewed stories before they were posted. But that would have been financial suicide. Digg's spectacular growth over the last few years was largely thanks to the site's unique story-selection technology. Converting Digg into a traditional news site would have alienated the vast majority of Digg's users and severely restricted the site's potential for continued growth.

The paradoxical result was that although Digg's executives had complete physical control over its servers, it faced significant constraints on its ability to control what appeared on its home page. Digg could remove any given story, ban any given user, and even install software filters that automatically removed content that matched certain programmati-cally defined criteria. Yet as long as Digg retained the user-generated editorial strategy that had been responsible for its success, Digg had no choice but to permit the stories its users wanted to put there. Ownership of the physical platform was no grant of the power to control its use.

Third-Party Instant-Messaging Clients

Another good example of the kinds of challenges a network owner can expect to face if it engages in wide-scale network discrimination can be found in the long-running battle over interoperability between the major instant-messaging networks owned by Microsoft, AOL, and Yahoo! and the developers of third-party instant-messaging applications. As the name suggests, instant messaging is a class of applications that allow users to see when their friends are online and chat with them in real time. Yahoo!, Microsoft, and AOL all offer clients free instant-messaging software and maintain servers that allow these clients to communicate.

A variety of third-party developers have developed competing instant-messaging applications designed to be compatible with these networks. Yahoo!, AOL, and Microsoft would prefer that users use their own client software instead. On several occasions in 2003 and 2004, Yahoo! made changes to its instant-messaging protocol designed to shut out third-party applications.[42] Microsoft and AOL made similar efforts to block third-party IM clients from their IM networks.[43] In each case, the developers of the third-party applications reacted swiftly, releasing software updates within a matter of days, if not hours, that re-established compatibility. Indeed, over time, the responses of the various third-party applications became more sophisticated and better-coordinated. They increasingly used shared libraries so that workarounds could be developed once and then deployed to all clients simultaneously.

Ultimately, AOL, Microsoft, and Yahoo! all relented. The constant software updates were a source of irritation for users of their official client software, and it became clear that users

Although Digg's executives had complete physical control over its servers, it faced significant constraints on its ability to control what appeared on its home page.

of third-party clients were not going to give up. Today, third-party clients such as Trillian, Pidgin, and Adium (and newcomers like Meebo) support all major instant-messaging networks with the grudging acquiescence of the networks' operators.

To be clear, Digg and instant messaging operate at the edges of the network, so these controversies are not network neutrality issues as such. But the stories suggest the kinds of problems that ISPs would be likely to encounter if they attempted to restrict their customers' use of their Internet connections. The TCP/IP protocols, like Digg's editorial process, place a great deal of power in the hands of end users. That power cannot be withdrawn by fiat. For almost a decade, broadband ISPs have accepted a basically passive role, transmitting the users' packets without interfering with them or even knowing what they contain. This decentralization of responsibility has made possible a breathtaking variety of interesting and useful applications, including Google, Wikipedia, World of Warcraft, AOL Instant Messenger, the iTunes store, and millions of others.

The assumptions of the open network are deeply embedded in each and every one of these applications. They are not designed for centralized control or administration. In the short term, network owners certainly have the power to block any given application, just as Digg has the power to delete any given story or AOL has the ability to block any given user from its IM network. But such blunt instruments aren't likely to succeed or to serve the network owner's interests in the long run. As we will see below, network owners' efforts to manipulate users' online activities are far more likely to generate ill will and spur the development of workarounds than they are to foster docile acceptance and higher profits.

Evasion for the Masses

Some advocates of government regulations requiring network neutrality are worried not that cutting-edge technologies would be blocked entirely but that network owners might dissuade enough nontechnical users to stunt the adoption of new, cutting-edge technologies. They worry that if only hard-core hackers can take advantage of workarounds, the effect will be little different than a total block.

But one lesson of the instant-messaging wars is that workarounds can often be quickly integrated into user-friendly consumer products that are accessible to ordinary users. Even the ragged band of volunteers and small start-ups that developed alternative instant-messaging applications early in this decade managed to quickly develop and distribute reasonably user-friendly applications that incorporated the latest workarounds. Users only needed to download an installer and double-click on it. While there are certainly some users who cannot manage this feat, there are tens of millions who can.

More important, a network owner launching a concerted assault on network neutrality would face much larger, better-funded, and more determined opponents. Major Internet firms such as Microsoft, Apple, and Google have a strong incentive to preserve the open Internet. It's not hard, for example, to imagine Google bundling circumvention software with the Google toolbar that's already on millions of Americans' computers. And, of course, Windows and Mac OS already have automatic software update features that could be used to distribute workarounds. These companies clearly understand that the erosion of the end-to-end principle could leave them at the mercy of network owners. They therefore have every incentive to design their products in ways that shift power toward themselves or their customers rather than to network owners.

"More Is Different"

Another challenge facing platform owners wishing to control their users' behavior is that platforms become progressively harder to control as they become larger and more complex. In his recent book, *Here Comes Everybody*,[44] Clay Shirky writes (quoting physicist Phillip Anderson) "more is different." That is, the behaviors of large, complex systems cannot easily be predicted from the behaviors of sim-

pler systems.

Should they try to implement new, more discriminatory policies, the owners of networks are likely to find that managing a complex platform with tens of millions of users is very different from managing a simple platform with tens of thousands or hundreds of thousands of users. As a network becomes larger and more complex, a larger, more hierarchical, and (inevitably) more bureaucratic organization will be needed to manage it. And the more control the organization attempts to exert, the more personnel it will require and the more acute the organizational challenges it will face.

We have already seen the difficulties faced by Digg and the various instant-messaging networks when they tried to limit their users' activities on those platforms. The Internet as a whole is an incomparably larger, more complex, and more decentralized system. If blocking unwanted news stories from Digg or unauthorized clients from AOL's instant-messaging network was difficult, blocking undesired content or applications from the Internet as a whole may be virtually impossible.

Of course, this isn't to say that the major network owners are completely powerless. In the short run, they can certainly block any given application or website. But randomly blocking a handful of websites or applications is unlikely to be a profitable business strategy. The discriminatory business models that network neutrality advocates fear require a sophisticated and comprehensive regime of price discrimination, and it is far from clear that it would be feasible to enforce such a scheme for the Internet as a whole.

There are millions of small websites, applications, and content providers. For any strategy of network discrimination to succeed, it would require, at a minimum, software that can identify and classify this heterogeneous traffic in real time. But the sheer number and variety of applications would make the development of such software extremely costly. In practice, the software would have to simply block any traffic it didn't recognize, which

would mean inconveniencing the millions of customers who use one or more uncommon applications.

As in the Digg incident, any effort by a network owner to exert more control over its portion of the Internet would face determined resistance from geek-activists who would develop creative ways to evade the filters. Just as Digg users transformed the AACS key into songs and pictures to evade text-based filtering, so hackers would develop software to camouflage disfavored traffic. The sheer number of potential adversaries would make organizing an effective response a monumental challenge.

In short, ISPs that attempt to limit their users' online activities are likely to learn the same lesson that Digg did: openness is a one-way ratchet. Once a firm cedes control to its users, things evolve in a way that makes it extremely difficult to reassert control. On a closed network, most users are unaware of the limitations being imposed on them, so they may not agitate for more openness. But once users have had a taste of freedom, they become acutely aware of any new restrictions and will stubbornly resist efforts to impose them.

Network Discrimination in the Real World

There is a common, but unstated, assumption in much writing about network neutrality that the Internet's open architecture is a fragile system that could collapse at the first sign of pressure. Advocates of new regulations point to scattered examples of network owners violating the end-to-end principle and suggest that these violations presage a more general retreat from a nondiscriminatory network.

But such pessimism is unwarranted. On a network with thousands of firms and hundreds of millions of users, it is not surprising that we see occasional deviations from the end-to-end principle. But such discriminatory policies have tended to be haphazard and rare. They have been minor headaches for a small number of broadband users rather than a

Once a firm cedes control to its users, things evolve in a way that makes it extremely difficult to reassert control.

threat to the Internet's fundamental architecture. And despite the recent bluster of some telecom executives,[45] efforts to undermine the end-to-end principle do not appear to be growing more frequent or more ambitious.

Indeed, the Internet has a rich history of being used in ways that were officially prohibited by the network's owners. This has been true since the early days of the ARPANET. In 1972, a single connection to the ARPANET could cost more than $100,000—half a million 2008 dollars.[46] Much of that cost was born by ARPA itself, and at least on paper, use of the network was to be restricted to ARPA-related research projects. Yet by the mid-1970s, there were unsanctioned mailing lists on the ARPANET devoted to weighty topics such as science fiction. Some such activities may have had the tacit approval of ARPA as a way of generating useful test traffic.[47] But when day-to-day operation of the network was transferred from ARPA to the Defense Communications Agency, military bureaucrats made a serious effort to crack down on "frivolous" uses of the network. For example, one 1982 message from DCA threatened to cut off sites that forwarded an "e-mail chain letter" that had been making the rounds.[48] However, those efforts had limited success. Thanks to the network's decentralized architecture, DCA's ability to monitor and control the use of its network was extremely limited.

In a 2003 paper, Tim Wu documented a variety of restrictions that broadband providers placed on their users earlier in this decade.[49] Wu's thesis was that these restrictions were a threat to open architecture of the Internet. But five years later, a different conclusion suggests itself: these restrictions, while irritating to individual customers who have been subject to them, have been too sporadically enforced to have had any real effect on the open character of the Internet.

For example, several cable providers prohibited customers from installing home networking equipment or sharing their Internet access with others outside of their premises. The current (as of July 2008) Comcast acceptable use policy includes a similar provision

prohibiting the use of WiFi to share Internet access with anyone outside of the customer's premises. Yet there's no way broadband providers could possibly enforce these restrictions in a systematic manner. Indeed, open WiFi networks have become quite common, and broadband providers do not appear to be taking action against their owners.

In addition, Comcast's current acceptable use policy does not allow users to "post, store, send, transmit, or disseminate any information or material which a reasonable person could deem to be indecent, pornographic, harassing, threatening, hateful, or intimidating." Although hard data on pornography dissemination by Comcast's customers is hard to come by, anecdotal evidence suggests that this restriction is not being enforced.

One hundred percent of cable providers and a third of DSL providers limited the operation of servers in 2002, a restriction that continues to appear in Comcast's latest acceptable use policy. And it has been at least sporadically enforced. However, there is little reason to think this restriction has been a significant obstacle to the development of innovative server software. First, there is a vibrant market for third-party hosting services, with prices as low as $10 per month.[50] This is well within the budget of anyone wanting to host his or her own content. Second, even those ISPs that ban the use of traditional servers generally permit (with some exceptions discussed below) the use of consumer applications, such as peer-to-peer applications, that have server-like characteristics. Most users do not want to run their own Web or e-mail servers, and would be unlikely to do so even if it were permitted. But when a significant number of users have begun to use applications that perform server-like roles, ISPs have generally not classified them as servers or attempted to restrict their use.

The Comcast Kerfuffle
Probably the most clear-cut example of a recent attack on network neutrality was last year's revelation that Comcast had been interfering with peer-to-peer file sharing traffic. In October, the Associated Press con-

Once users have had a taste of freedom, they become acutely aware of any new restrictions and will stubbornly resist efforts to impose them.

firmed rumors that had been circulating on the Internet for months that Comcast was actively interfering with its customers' use of BitTorrent and similar peer-to-peer file sharing applications. The AP reported, and others subsequently confirmed, that Comcast's network would sometimes send forged "reset" packets to both ends of a peer-to-peer connection, effectively telling each end of the connection that the other had hung up.[51] Comcast is reportedly using software manufactured by a company called Sandvine to perform this feat.[52]

Comcast's defenders argued that this policy is necessary to combat congestion on its network. Those arguments weren't totally implausible. Peer-to-peer activity constitutes a large fraction of online traffic, and the networking technology currently in use by Comcast is designed for fast downloads at the cost of slow uploads. Comcast argues that peer-to-peer traffic can place unique stresses on its asymmetrical network.

Comcast's blocking didn't just affect the heaviest BitTorrent users who were downloading gigabytes of illicit movies and music. Even BitTorrent users engaged in totally innocuous (and relatively low-bandwidth) activities like downloading the latest bug fixes for the online game World of Warcraft could be affected.[53] And in an apparent misconfiguration, Sandvine also appears to have interfered with the popular Lotus Notes business software.[54]

Comcast's activities attracted considerable public attention. Comcast scrambled to explain its actions and insist (somewhat misleadingly) that customers will continue to "enjoy unfettered access to all the content, services, and applications that the Internet has to offer."[55] Comcast's competitors have relished the opportunity to tout their own, less discriminatory, network policies. Verizon, whose recent investments in fiber-optic lines give it significantly more bandwidth than Comcast, has crowed that its "more robust" network makes such filtering unnecessary. Verizon also took a shot at Comcast's secretive policies by pledging to let customers know before it filters

traffic in the future.[56]

But the most important development was the reaction of BitTorrent users themselves. In early 2006, BitTorrent developers began adding encryption features to BitTorrent clients to defeat traffic-shaping tools.[57] Within days of the AP story, BitTorrent users began swapping tips for evading Comcast's blocks. Most BitTorrent software supports a technique called "header encryption" that makes BitTorrent packets difficult for filtering software to identify.[58] The primary long-run effect of Comcast's interference with BitTorrent traffic won't be a reduction in that traffic, but simply more rapid adoption of encrypted versions of the BitTorrent protocol.

For a variety of reasons, including user resistance, negative publicity, and regulatory pressure, Comcast backed away from its discriminatory policy in March. It reached an agreement with BitTorrent, Inc., the company founded by BitTorrent creator Bram Cohen, to stop interfering with BitTorrent traffic. In return, BitTorrent agreed to work with Comcast to make the BitTorrent protocol "more efficient." Comcast pledged to implement an end-to-end-friendly traffic-shaping regime by the end of 2008.[59] By the time the FCC released a ruling on Comcast's behavior in July, the issue had already been rendered a moot point by technological and market developments.[60]

For all of its technical sophistication, Sandvine is still a relatively blunt instrument. No doubt some of its specific flaws will be fixed. But no amount of tinkering with a tool like Sandvine could give Comcast the kind of comprehensive control over its network that network neutrality advocates have warned about. Sandvine requires Comcast to describe the types of traffic it wishes to block in specific, technical terms. Yet the types of control network neutrality advocates hope to trump with regulation—"don't waste bandwidth," "don't share copyrighted files," or "don't use content or applications that compete with our affiliates"—are not based on technical criteria. They are business criteria that would require constant tweaking by an army of network engi-

Discriminatory policies have tended to be haphazard and rare.

neers to implement and maintain. Regulation is not needed to frustrate such controls. The open architecture of the Internet is sufficient.

Interfering with BitTorrent may have given some short-term relief to Comcast's aging network while the company upgrades to a new, higher-speed networking technology called DOCSIS 3.0.[61] But as Comcast evidently realized, it would not have worked as more than a stop-gap strategy. Over time, the vast majority of peer-to-peer users would have either learned how to evade Comcast's filters or shifted to competing firms such as Verizon that offer unfiltered Internet access.

At a minimum, it's clear that Comcast's practical ability to control its users' online activities is sharply constrained by technical and economic forces. Comcast's ability to exert fine-grained control over its users is much more limited than some network neutrality advocates fear.

Assessing the Threat to End-to-End

One of the challenges of evaluating the case for network neutrality regulation is that every supporter of new regulation seems to have a different idea about the types of discrimination that network owners are most likely to undertake. Some analysts suggest that network owners will focus narrowly on degrading applications, such as Internet telephony, that compete with their legacy businesses. Others suggest that they will undertake a broad scheme of price discrimination in which virtually every online application would face a choice between degraded service and higher fees. Still others are worried about risks of censorship.

One of the clearest statements of the scenario network neutrality advocates are concerned about was voiced by Ed Whitacre, then the CEO of SBC (which soon became AT&T). In a 2005 interview with *Business Week*, Whitacre created a firestorm of controversy when he argued that large Internet firms like Google, Vonage, and MSN should pay his company for the privilege of reaching SBC's customers. In an unusually candid moment, Whitacre stated "what they would like to do is use my pipes free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it. So there's going to have to be some mechanism for [firms like Microsoft, Google, and Vonage] who use these pipes to pay for the portion they're using."[62]

One model of this kind of network discrimination was fleshed out by Harvard law professor Yochai Benkler in his widely cited 2006 book, *The Wealth of Networks*,[63] which examines the social, economic, and political consequences of the growth of the Internet and digital technologies more generally. He argues that private ownership of communications infrastructure—even in a competitive market—threatens the open character of the Internet, which in turn threatens both innovation and free speech.[64]

Figure 1 illustrates Benkler's simple model for a world without network neutrality regulation. In this model, D might be a major Internet service provider, A might be a residential broadband customer, and B and C might be competing website operators. Benkler argues that this setup gives D control over A's "information environment," raising three fundamental concerns about this arrangement: First, D is in a position to charge B and C unreasonable prices to communicate with A. Second, D is in a position to sign an exclusive contract with B guaranteeing that B's messages reach A but some or all of C's do not. Finally, D may be able to alter messages to or from A as they pass through D's network, thereby misleading or manipulating A for the benefit of D or another party's benefit.[65]

In this simplified model, it seems self-evident that D has almost unlimited power over A's online experience. If B or C wants to transmit a message that D feels is contrary to its interests, D can fail to deliver the message or even alter the message to suit its purposes.

No doubt, Benkler would acknowledge that this is a greatly simplified model for the real-world Internet. He recognizes that D's ability to manipulate A is limited by "the degree to which it is hard or easy to get around

By the time the FCC released a ruling on Comcast's behavior in July, the issue had already been rendered moot by technological and market developments.

Figure 1
A Simple Model of Communications Networks Inspired by Yochai Benkler

Source: Cato Institute.

D's facility," and he also argues that the "the degree of transparency" of D's manipulations is important.[66]

But Benkler nevertheless underestimates the ability of users to detect when a network owner is manipulating their traffic. He writes that

> there are many reasons that different sites load at different speeds, or even fail to load altogether. Users, the vast majority of whom are unaware that the provider could, if it chose, regulate the flow of information to them, will assume that it is the target site that is failing, not that their own service provider is manipulating what they can see.[67]

As we have just seen, Comcast's actions were relatively subtle and narrowly targeted BitTorrent, a relatively obscure and unpopular protocol. Yet it took only a handful of tech-savvy users to pinpoint which carriers were blocking traffic and how they were doing it, and that information was picked up by the mainstream media and widely publicized.

Benkler also suggests that D's control over A's "information environment" would be so complete that D could prevent nontechnical users like A from even learning about D's manipulation. But in the real world, even a broadband monopolist won't have a monopoly on the spread of information. People have dozens of information sources, including broadcast television and radio, newspapers, magazines, cable and satellite television, telephones, and face-to-face conversation. In the real world, D's ability to control A's "information environment" is destroyed by the multiplicity of alternative information sources.

Theoretical models can be useful when they remove extraneous details and focus the reader's attention on the essential characteristics of a problem. And at first glance, it seems obvious that party D in Figure 1 (say, AT&T)

Comcast's practical ability to control its users' online activities is sharply constrained by technical and economic forces.

19

Figure 2
A More Realistic Model of the Internet

Source: Cato Institute.

can charge party B (Google, Microsoft, or a Web startup) for access to A (a broadband customer). But the simplifications of Benkler's model obscure more than they illuminate.

A more realistic model of the Internet is presented in Figure 2. Two new parties, E and F, have been added, and B and C have been moved so that they now receive their connection to D (and, indirectly, to A) through them. Here, D might represent a broadband provider like AT&T, and E and F might represent competitive backbone providers like Global Crossing or Level 3. At first glance, this might not seem to change D's bargaining position at all. D is still the exclusive gatekeeper for traffic to A. D can still selectively block messages from B to A, and D can still accept payments from B to block messages from C to A.

However, the picture has changed in sev-

eral important ways. As noted previously, more is often different. If F had only a handful of customers, it would be feasible for D to approach each of them and credibly threaten to disconnect them from D's customers if they didn't pay up. But as the number and variety of F's customers grows, logistical concerns become more and more significant. D doesn't have a direct contractual relationship to C, and C's traffic to A is intermingled with the traffic of F's other customers. F has no particular reason to inform D when C joins the network, nor to furnish D with contact or billing information for C.

To illustrate, this author contributes to the Technology Liberation Front,[68] a small group blog that is administered by a friend who runs a small web-design business. The site generates a negligible amount of advertising revenue and has no employees. The amount of revenue a major backbone provider could extract from the blog would be too small to even recoup the costs of classifying the site, identifying the site's owner, obtaining contact information, discussing the payment options with the site's owner, and so forth. Not only would it be a waste of money for a network owner to try to extort money from the blog, the site's contributors would respond by publicly ridiculing that network provider.

This small example would be multiplied hundreds of thousands of times if a network owner were to approach sites with a wide variety of business models, organizational structures, and financial situations for payment. A lot of sites, *especially* the small ones, would credibly refuse to pay a dime.

That would create a chicken-and-egg problem. To return to Figure 2, D's primary revenue source is A and A's cohorts, who pay D for the ability to communicate with C and other customers of F. D cannot block all of F's customers (who might collectively represent a significant fraction of all content on the Web) from reaching A without sparking cancellations—if not lawsuits—from its own customers. But C won't be inclined to pay D unless C believes D's threat to cut off C is credible. C knows that D's threat won't become credible until a significant number of F's other customers have paid the fees D has demanded. So although D may make a lot of threats, it will be relatively obvious that D can never actually pull the trigger and cut its own customers off from thousands of websites.

The graph in Figure 2 has an important symmetry. D, E, and F are "Tier 1" backbone providers. The Internet has approximately nine "Tier 1" network owners that peer with one another on a settlement-free basis.[69] Smaller network owners pay one or more of these carriers to carry at least some of their traffic. A randomly selected pair of Internet users will most often be connected via different backbone providers, and the path between them will therefore traverse a settlement-free peering point. So it is reasonable to take this as the typical case.

D needs access to F's customers roughly as much as F needs access to D's customers. Moreover, F has a strong incentive to thwart any effort by D to charge its own customers for interconnection, because that would put F at a serious competitive disadvantage in the competition for the business of customers like B and C. F will be strongly motivated to organize a "common front" with E against D. If necessary, F is likely to provide legal, public relations, and perhaps even direct financial support to any customers that are targeted by D in order to ensure that D's efforts fail.

In the face of well-organized opposition, it is extremely unlikely that D would be able to extract payments from enough of F's customers to make credible a threat to cut off the rest. And that, in turn, means that C will have no particular incentive to comply with D's demands. This is especially true because C would know that if it agreed to pay D for access to D's customers, numerous other ISPs would demand similar terms for themselves. From C's perspective, being temporarily cut off from D's network would be far preferable to signaling to the world that it was susceptible to such bullying.

In the event of a prolonged standoff, in which D blocked C from access to A, it is almost certain that D would blink first. Such a move

Even a broadband monopolist won't have a monopoly on the spread of information.

would generate enormous amounts of positive publicity for C and negative publicity for D. It would likely spark customer cancellations, and perhaps even lawsuits based on false advertising, breach of contract, or the like. A prolonged standoff would do tremendous damage to D's reputation and considerable damage to its bottom line. Indeed, if C were significantly smaller than D, C might even *relish* the opportunity to be cast in a David-and-Goliath battle with D. Anyone who got their Internet access from someone other than D would still be able to reach C, and it's likely that C's total traffic would grow as thousands of people came to learn more about the site D was so determined to block.

Whitacre's successors at AT&T would no doubt love to charge the Googles and eBays of the world for the privilege of accessing AT&T's broadband customers. But the structure of the Internet simply doesn't give AT&T the leverage to do so. AT&T needs Google and eBay (and millions of smaller websites) as much as they need it.

Exclusive Access

Another way Benkler suggests D might profit from its exclusive connection to A is by accepting payments from B to limit communication between A and C. In this case, B and C might be competitors, and A might be an existing customer of C's online service. B might calculate that it could pay D to cut off (or slow down) A's connection to C and thereby win A's business.

Just stating the proposal in those terms makes the first difficulty clear: A is going to wonder why he is suddenly unable to reach C, and once the media report that it's because B paid D to cut off access, he's not going to be enthusiastic about becoming a customer of B. Indeed, it's likely that such a move would lead a significant number of B's customers who are *not* connected through D to switch to C in protest. As in the previous case, becoming the target of a deliberate block by D would be a tremendous publicity coup for C.

Moreover, B doesn't simply want to block access to C. B probably wants to block access to anyone offering a site that competes with B. If, for example, B is an online video site, B will want to ask D to block all competing online video sites. Initially, that will be easy: B can just provide D with a list of the major competing video sites, and D can institute a block based on the relevant IP addresses. But preserving B's monopoly on online video would grow more difficult over time. B's competitors are likely to begin looking for ways to route around D's block. B or D would need to expend considerable resources to track and counter those efforts and ensure that competing video services remain inaccessible. Not only would preventing the emergence of new video applications likely be far more trouble that it's worth in the long run, but each time a new site got blocked it would create a new wave of negative publicity highlighting the inadequacies of B's product.

Free Speech

Benkler also worries that by interfering with a customer's communications, a network provider would be able to manipulate that customer's "information environment" and thereby skew the user's view of the world.[70] For example, a conservative network owner might arrange things so that it would be easier for its customers to reach conservative-leaning blogs and news sites and harder for customers to reach liberal-leaning ones. However, Benkler never explains in any detail how the network owner would accomplish such a feat, and the Internet has strong built-in safeguards against network owners manipulating online content without the user's knowledge. The fundamental safeguard is the vast amount of content available on the Internet. Manipulating human communications effectively can only be performed by other human beings, and no company in the world has enough manpower to review every blog post, YouTube video, podcast, and e-mail on the Internet. Even evaluating every website in order to create a content-based blacklist would be a Herculean task.

Moreover, even if a company had the manpower required to evaluate all the content on the Internet, it's not at all obvious what techniques would be available to manipulate customers.

A randomly selected pair of Internet users will most often be connected via different backbone providers, and the path between them will therefore traverse a settlement-free peering point.

Benkler suggests that network owners might "make some sites and statements easier to reach and see—more prominently displayed on the screen, faster to load."[71] But that wouldn't make a lot of sense. Slowing down a user's access to Paul Krugman is unlikely to cause the user to read Michelle Malkin instead. And there's no obvious way to gracefully yet surreptitiously make a website "more prominently displayed on the screen," since users directly control which websites they visit. Users would surely notice if their ISP somehow began causing unwanted websites to pop up on their screens.

Even if all of these technical and logistical hurdles could be overcome, there remains a fundamental problem of backlash. Benkler suggests that such techniques would be "subtle," but they would in fact be nothing of the sort. No widescale manipulation would go unnoticed for very long. There are a variety of network diagnostic tools that can be used to analyze response times from various parts of the Internet and determine the source of the interference.[72]

As we saw with the Digg/AACS incident, when a powerful individual or organization tries to suppress speech it dislikes, it tends to trigger what Techdirt blogger Mike Masnick has dubbed the "Streisand Effect," named after an aerial photo of actress Barbara Streisand's house that became famous after the media extensively covered her lawsuit seeking to have it removed from the Internet.[73] In 2006, dozens of high-traffic blogs reported on allegations that Kentucky state employees had been blocked from viewing a liberal-leaning blog on state computers.[74] The outrage over that relatively petty censorship pales in comparison to the firestorm of controversy that would be unleashed if a major network owner embarked on a systematic campaign of censorship on its network.

We don't have to speculate on what such a backlash would look like. We have only to look at the controversy over Verizon Wireless's decision to deny a pro-choice group access to an SMS "short code," a number used to send and receive text messages. The decision attracted near-universal condemnation that forced Verizon to back down a week later.[75] This was not a true network neutrality issue, but it's a good illustration of the potency of public opinion when large companies are perceived to be interfering with free speech. The backlash against a company engaging in deliberate, widescale censorship on the Internet would be even more severe.

The Role of Competition

It's worth noting that none of the arguments in the preceding sections require a significant amount of competition in the residential broadband industry. Most of the difficulties that network owners would encounter if they deviated from the end-to-end principle would exist even if they had a monopoly of Internet access. Monopolists generally seek to maximize profits. Discriminatory strategies that reduce the value of the network without generating significant revenues are going to be money-losers whether or not the network owner holds a last-mile monopoly.

Of course, these arguments are even stronger in places where there is a broadband "duopoly." And they will be stronger still if new technologies—for example, broadband over power lines, WiMax, or higher-speed Internet access via the recently auctioned 700 MHz spectrum—introduce a third or fourth broadband option to a significant number of consumers. But even in the current market environment of relatively limited competition, broadband providers would still find it difficult to undermine the end-to-end principle.

Some of the arguments above *do* assume a robust and competitive market for wholesale access to the Internet backbone. Luckily, there is ample competition in this market, with competitive "tier one" backbone providers such as Level 3 and Global Crossing competing on a roughly level playing field with the largest "last mile" broadband providers.

End-to-End and the Incentive to Innovate

Hundreds of Internet-based startups are founded each year in the hopes that they will be the next Yahoo!, eBay, or Google. Startup founders work long hours and take great per-

AT&T would love to charge the Googles and eBays of the world for the privilege of accessing AT&T's customers, but the structure of the Internet simply doesn't give AT&T the leverage to do so.

sonal risk to build their firms. Many network neutrality advocates worry that Internet discrimination could deter the creation of new startups and slow the pace of high-tech innovation.[76] As mentioned previously, the Internet's open architecture allows firms to enter the market without seeking the approval of the hundreds of companies that control the various networks that collectively make up the Internet. If the Internet were transformed into a proprietary network, that would certainly slow the creation of online startup firms.

But it is overstating the case to suggest that even sporadic interference with the end-to-end principle, such as Comcast's interference with BitTorrent, significantly reduces the incentives for online innovation. The sheer number of startup firms gives them a kind of "safety in numbers." The typical startup's odds of being targeted by a major network provider are quite low. Moreover, the fact that dozens of different companies own significant parts of the Internet's infrastructure means that even those firms unlucky enough to be targeted by one network provider will still be able to reach the vast majority of Internet users via other networks. To be sure, such discrimination would be a headache for these firms, but a relatively small chance of being cut off from a minority of residential customers is unlikely to rank very high on an entrepreneur's list of worries.

There is, in short, little reason to believe that network owners will find it profitable to block or degrade content or applications on their networks. Openness is a one-way ratchet, and the end-to-end principle has developed sufficient inertia over the last three decades that it will be extremely difficult for network owners to displace it.

The Fast Lane and the Walled Garden

Thus far, we have considered scenarios in which incumbent broadband providers could intentionally degrade the performance of disfavored applications or content, and we have seen that such degradation is unlikely to

be profitable for network owners. We now consider the flipside of this strategy: an ISP that maintains a baseline level of connectivity for all applications and content, but selectively provides enhanced connectivity for applications or content owned with the network owner or its partners.

In this "fast lane" scenario, colorfully described in dozens of popular accounts over the last two years, the Internet would be divided up into two "lanes."[77] The Googles and Microsofts of the world would enjoy speedy delivery in the "fast lane." Everyone else's traffic would be relegated to the "slow lane." Activists worry that this will stifle innovation, as only the select few who can pay the freight on the "fast lane" will be able to deliver next-generation services, while everyone else's online offerings stagnate.

Such a strategy would not pose a short-term threat to the Internet's end-to-end architecture. Existing applications would enjoy the same level of bandwidth they had always enjoyed and would continue to operate normally. New applications that required only today's level of network performance would be able to launch without difficulty. That, in turn, means that if a "fast lane" strategy required regulatory intervention, policymakers would have plenty of time to study the problem and craft a response after the fact.

The fundamental difficulty of the "fast lane" strategy is that a network owner pursuing such a strategy would be effectively foregoing the enormous value of the unfiltered content and applications that comes "for free" with unfiltered Internet access. The unfiltered Internet already offers a breathtaking variety of innovative content and applications, and there is every reason to expect things to get even better as the available bandwidth continues to increase. Those ISPs that continue to provide their users with faster, unfiltered access to the Internet will be able to offer all of this content to their customers, enhancing the value of their pipe at no additional cost to themselves.

In contrast, ISPs that choose not to upgrade their customers' Internet access but instead devote more bandwidth to a proprietary "walled

garden" of affiliated content and applications will have to actively recruit each application or content provider that participates in the "fast lane" program. In fact, this is precisely the strategy that AOL undertook during the 1990s. AOL was initially a proprietary online service, charged by the hour, that allowed its users to access AOL-affiliated online content. Over time, AOL gradually made it easier for its customers to access content on the Internet, so that by the end of the 1990s, it was viewed primarily as an Internet Service Provider that happened to offer some proprietary applications and content as well.[78] The fundamental problem requiring AOL to change was that the content available on the Internet grew so rapidly that AOL (and other proprietary services like CompuServe) couldn't keep up. AOL finally threw in the towel in 2006, announcing that the proprietary services that had once formed the core of its online offerings would become just another ad-supported website.[79] A "walled garden/slow lane" strategy has already proven unprofitable in the marketplace. Regulations prohibiting such a business model would be surplusage.

Large, hierarchical organizations face great difficulties keeping up with the innovation of a decentralized, open platform. The inherent frictions in managing and expanding a proprietary online service makes it virtually impossible for the owner of a "walled garden" to innovate as rapidly as thousands of companies competing on an open platform. Even a firm as large and well-capitalized as AT&T, Verizon, or Comcast will have difficulty developing a stable of content and applications that will be as appealing as the content and applications available on the unfettered Internet.

Indeed, as Chris Yoo has pointed out,[80] the merger of AOL and Time Warner in 2001 was conceived as just such a vertically integrated network/content juggernaut. The merged firm proved to be an anemic competitor. As impressive as Time Warner's stable of content was, it was dwarfed by the content already available on the open Internet. AOL's dial-up division simply could not have afforded to cut off its customers' access to unaffiliated Internet content, because doing so would have dramatically reduced the value of its online offering.

Of course, in some areas, consumers may not have the option of purchasing unfettered access to the Internet, either because they have only one broadband provider, or because both broadband providers in their area are pursuing "walled garden" strategies. But even in those areas, several factors will create pressure on ISPs to provide full-speed Internet access alongside their "walled garden" services. First, even a monopolist has an incentive to maximize the monopoly rents he can extract. If a high-speed connection to the unfiltered Internet is significantly more valuable to customers than access to the walled garden, then offering unfiltered Internet access will be a revenue-maximizing strategy even in the absence of competition. Indeed, the incentive is likely to grow over time, as the content on the unfiltered Internet gets further and further ahead of what is available within the "walled garden."

This is even more true in a "duopoly" situation, in which an area is served by both a cable and a phone incumbent, and both are pursuing a "walled garden" approach. Each firm would have a powerful incentive to "break ranks" and increase the speed of their unfiltered Internet access, thereby attracting a significant number of customers from the other carrier. Moreover, because of the patchwork nature of the cable and phone companies' service areas, almost every large incumbent phone company has several cable competitors in parts of its service territory, and vice versa. For a variety of practical reasons, ISPs are unlikely to offer unfiltered Internet access to some of its customers and limit other customers to a "walled garden," so as long as they face competition from the unfiltered Internet in a significant number of markets, they will have good reasons to continue offering it across their service areas.

High-Definition Video

The "walled garden" strategy is almost always described in terms of high-definition video, because that's currently the most important application for which existing Internet connection speeds are inadequate. Those who fear a "two-tiered Internet" worry that carriers will find

ISPs that continue to provide their users with faster, unfiltered access to the Internet will enhance the value of their pipe at no additional cost to themselves.

it more profitable to devote most of their bandwidth to a handful of large media companies—whose content can be sold at a significant markup—than to allow consumers to use that bandwidth to freely access any content on the Internet.

The problem with this theory is that many broadband networks *already have* a proprietary "fast lane" for video: cable television. Cable firms have always allocated the bulk of the bandwidth on their coaxial cables to video transmission, not Internet access. Recently, Verizon and AT&T have been rolling out proprietary video services of their own. While originally, cable television was an analog service that was dramatically different from modern data networks, cable operators have been steadily moving to Internet-based technologies. AT&T's U-Verse video service, for example, will reportedly be based on TCP/IP networking technologies.

We have a lot of experience with the economics of proprietary video networks, and there is no reason to think that broadband firms could generate significant revenues by selling "fast lane" access to HD-video producers. Indeed, on cable television networks, the money flows the other way, with the network owner paying the content provider for the privilege of carrying its content.[81] There's nothing about the transition to TCP/IP-based content delivery that would strengthen the network owners' bargaining position enough to cause the payments to begin flowing in the opposite direction.

Of course, partisans for open networks would prefer that 100 percent of the available bandwidth be allocated to unfiltered Internet access. But cable and telephone incumbents have already invested billions of dollars in video-on-demand infrastructure. Legislation requiring all cable and telephone bandwidth to be reallocated to public Internet connectivity would be a nonstarter politically as a transparent seizure of private investment. The leading network neutrality proposal of 2006, known as Snowe-Dorgan,[82] explicitly exempted cable television services from network neutrality requirements. Legislation requiring data networks to be nondiscriminatory will have no impact on the amount of bandwidth devoted to proprietary—and network-neutrality-exempt—video services.

It would be overstating the case to suggest that no ISP will attempt a "walled garden" strategy and construct a "fast lane" to promote its success. But the Internet is not so fragile that a few "walled gardens" pose a threat to its vitality. The Internet is much bigger than any one network owner, and under any conceivable scenario, there will continue to be hundreds of millions of people with unfettered, high-speed access to the open Internet. Walled gardens are likely to prove anemic, unprofitable, and (as a consequence) short-lived. And because a "fast lane" strategy is unlikely to interfere with existing Internet applications, policymakers can afford to wait until any problem manifests itself before taking action.

To summarize, network neutrality supporters have suggested two basic ways that network owners might profit from undermining the end-to-end principle. One strategy involves threatening to degrade or block applications or content as a way of getting Internet firms to pay extra for unfettered access to their customers. The other strategy involves selling access to a proprietary "fast lane" that gives preferential treatment to affiliated applications or content. While it's not inconceivable that network owners will try either or both of these strategies, they are likely to prove unprofitable and as a result will be short-lived.

Government regulation to protect the Internet's end-to-end architecture is unnecessary because a variety of nonregulatory forces are sufficient to prevent it. But if network neutrality regulation were merely unnecessary, it might make sense to enact it anyway just to be on the safe side. History suggests that regulation is likely to prove not only unnecessary but harmful as well.

The Deregulation Revolution

In the first half of the 20th century, the study of government regulation was dominated by a collection of ideas that came to be

If a high-speed connection to the unfiltered Internet is more valuable to customers, then offering unfiltered Internet access will be a revenue-maximizing strategy even in the absence of competition.

known as the "public interest" theory of regulation. Under the influence of Progressive theories of political economy, New Deal–era economists and policymakers had great confidence in the ability of expert government regulators to correct perceived market failures through active intervention in market processes.[83]

They created or strengthened numerous regulatory agencies; three of the most important are the agencies that collectively oversaw the nation's transportation and communications infrastructure: the Interstate Commerce Commission, the Federal Communications Commission, and the Civil Aeronautics Board. But then in the 1960s and 1970s—at the same time that computer scientists were developing the infrastructure and ideas that would power the Internet—another intellectual revolution occurred in the public policy world. The result was a dramatic deregulation that has had profound effects on the American economy.

A bit of history about these agencies is crucial to understanding the intellectual revolution of the 1970s and its implications for today's regulatory debates.

The Interstate Commerce Commission

The Interstate Commerce Commission emerged from a debate strikingly similar to today's network neutrality debate. In the 1880s, the railroads were a new, vibrant industry in the process of transforming the Ameri-can economy. Activists became alarmed at their rapidly increasing size and power.

Congress responded in 1887 with the Interstate Commerce Act, which created the first modern regulatory agency, the Interstate Commerce Commission. Using language strikingly similar to modern network neutrality proposals, the ICA prohibited the railroads from charging different rates for "like and contemporaneous service in the transportation of a like kind of traffic under substantially similar circumstances and conditions." It also prohibited giving "undue or unreasonable preference or advantage" to any particular customers. Complaints regarding violations of these rules could be directed to the ICC or directly to the courts.

This was intended to protect consumers and the public from powerful companies.

But the man President Cleveland chose as the first ICC chairman, Thomas M. Cooley, was a railroad ally,[84] and ICC regulation of the railroads was relatively weak under his tenure.[85] The ICC was rendered even more impotent by a wave of litigation that engulfed the commission in the 1890s. Courts began second-guessing the rates the ICC tried to impose on the railroads, and Cooley began to worry that the railroads would ignore the ICC's decisions entirely.[86] The commission reached a low point in 1897 when the Supreme Court denied that it had been given the power to set rates at all.[87] The ICC was regarded as basically toothless for the next few years.[88]

Congress beefed up the commission's authority in 1903, 1906, and 1910.[89] But even after doubts about its legal authority were laid to rest, the ICC pursued a policy of general timidity, leaving in place discriminatory rate-making policies that had become long established and that had the support of powerful interest groups. The commission did reject a few rate increases in the years leading up to World War I, but in general, the ICC used "its considerable strength to preserve the status quo."[90]

Things got much worse after the war, as federal railroad regulation took on an overtly protectionist cast. In 1920, any pretense of protecting consumers was dropped, as Congress passed legislation giving the ICC the power to establish minimum as well as maximum rates.[91] In 1935, Congress reacted to "cutthroat" competition from truckers by extending the commission's authority to that industry as well.[92] ICC authority was extended to water shipping in 1940.[93]

For the next 40 years, the commission effectively operated a cartel for the benefit of transportation interests. In 1970, a report released by a Ralph Nader group described the commission as "primarily a forum at which transportation interests divide up the national transportation market."[94] Not only were consumers harmed by unnecessarily high prices, but economic efficiency was undermined because the ICC micromanaged the firms' activities, dictating which

Government regulation to protect the Internet's end-to-end architecture is unnecessary because a variety of nonregulatory forces are sufficient to prevent it.

routes they could serve and what cargo they could carry, meaning that competition could not drive down price and drive up quality. In many cases, trucks would carry cargo to a destination and then return empty because they were unable to secure permission from the ICC to carry cargo on the return trip.

The Civil Aeronautics Board

The ICC approach to regulation was at its zenith during the New Deal, and the commission served as a model for other regulatory bodies created during the period. One example is the Civil Aeronautics Board, which governed commercial aviation, taking a protectionist stance almost from its creation in 1938. In 1941, "the board first enunciated what was to be its philosophy on new entrants: the present number of carriers in air transportation was deemed sufficient to protect against monopoly, and any future expansion of air transportation would be best accomplished by the certification of presently operating air carriers."[95] In short, the CAB worked to *exclude* new competitors.

For the next 40 years, the CAB regulated the airline industry much the way the ICC regulated surface transportation. Both agencies were charged with the seemingly impossible task of simultaneously promoting the interests of consumers and incumbent firms. More often than not, consumers lost, as regulated firms cultivated cozy relationships with the regulators and used their influence to limit competition and raise prices.

The Federal Communications Commission

A more complicated case is the story of the telephone industry, but there too regulation ultimately worked as a barrier to competition. The FCC's efforts to protect telecom incumbents from competition took decades to break down.

Scholars disagree about the extent to which government regulation contributed to AT&T's initial dominance of the telephony market,[96] but all agree that after the nationalization of the telephone network during World War I, AT&T had a de facto monopoly on telephone service in the United States. This monopoly was formalized with the passage of the 1934 Communications Act, which put the newly created Federal Communications Commission in charge of regulating the Bell system.[97]

In 1942, the FCC responded to what it regarded as AT&T's excessive long-distance profits by requiring AT&T's long-distance operation to make payments to the Bell subsidiaries that provided local service.[98] Over the next three decades, the FCC required long-distance customers to bear an increasing share of the costs of local telephone infrastructure, effectively forcing long-distance customers to subsidize the cost of basic phone service.[99]

The growing gap between the cost of providing long-distance service and the prices AT&T charged to consumers created a large profit opportunity for any firm that could provide competitive long-distance service. To ward off this danger, the FCC strictly regulated entry into the long-distance market between 1942 and 1969. When new wireless communications technologies were developed that could have offered new competition, the FCC dragged its feet on approving their use. It approved the use of wireless microwave links for private lines (i.e., lines owned and used by a single firm) in the 1959 *Above 890* decision, but insisted that AT&T maintain its monopoly in offering long-distance service to the general public.[100]

In 1963, a startup firm called Microwave Communications Inc. (MCI) applied for permission to build a microwave link between Chicago and Saint Louis and lease access to other companies, which would make it a direct competitor to AT&T's long distance business. After six years of foot-dragging, and in a political climate that was beginning to favor competition over monopoly, the FCC finally approved MCI's application. In the 1970s, MCI applied for permission to build hundreds more links, creating what became the first competitive long-distance firm.[101]

The FCC's foot-dragging most likely delayed the introduction of long-distance competition by a decade or more. MCI had to wait seven years to get approval for its initial link between

New Deal–era economists and policymakers had great confidence in the ability of expert government regulators to correct perceived market failures.

Chicago and Saint Louis, and several more years after that before it could offer service to a significant fraction of the country. Like the ICC and the CAB, the FCC protected a client industry from the vagaries of markets and competition. As a result, they foreclosed new innovation and improvements in consumer welfare.

Deregulation

The deficiencies of these alphabet-soup agencies inspired a growing body of literature in the 1960s and 1970s that questioned the optimistic theories of regulation that had prompted their creation.[102] A new generation of economists began to realize that real-world regulators cannot be assumed to be either selfless or omniscient. Because regulators possess limited information, they are prone to many of the same types of errors that can degrade the performance of markets. Moreover, while there are certainly honest and public-spirited regulators, these economists questioned the common assumption that regulators will consistently be selfless advocates for the interests of the public.[103] The "revolving door" between regulators and the firms they regulate is eloquent testimony to the potential for corruption of the regulatory process.

Economists began to articulate theories of "regulatory capture," in which regulators increasingly come to serve not the public but organized interests with business before them.[104] They had little trouble finding examples, with the ICC, CAB, and FCC playing starring roles.

These academic critiques of regulation began to percolate in Washington, D.C. An early leader on the issue was Sen. Edward M. Kennedy (D-MA), who chaired the Senate Subcommittee on Administrative Practice and Procedure of the Judiciary Committee. He hired future Supreme Court Justice Stephen Breyer, then a Harvard law professor and an advocate of deregulation, to advise the committee.

The absurdity of the CAB's oversight of the airline industry made it an ideal poster child, and Breyer organized a series of hearings in 1975 to highlight the CAB's anticompetitive policies.[105] When President Carter took office in 1977, he appointed respected economist

Alfred Kahn to head the CAB. In a virtually unprecedented development, Kahn began dismantling his own agency's regulatory authority, steadily increasing the airlines' freedom to choose routes and set prices. Congress followed his lead in 1978, passing legislation that ratified the liberalization he had already undertaken and authorizing further movement in the same direction.[106] The result was increased competition and greater airline efficiency. From 1976 to 1982 alone, real fares fell by 8.5 percent despite fuel cost increases of 73 percent over the same period.[107]

The trucking industry was deregulated in a similar fashion. The ICC acted on its own authority to increase competition in the late 1970s, and many of these reforms were reinforced by Congress with the Motor Carrier Act of 1980. Deregulation dramatically increased competition and improved economic efficiency by reducing the number of trucks that made empty return trips due to regulatory restrictions, for example.[108]

Telecommunications is a more complicated case. Full-scale deregulation of telecommunications has yet to occur, but the trend toward deregulation did reach the industry in the late 1970s. The FCC continued to loosen restrictions on long-distance service throughout the 1970s, and Congress rebuffed AT&T's efforts to secure legislation reestablishing its long-distance monopoly. Then, under the combined pressure of the FCC, Congress, and a Justice Department antitrust complaint, AT&T agreed in 1982 to divest its local operating companies, creating a competitive long-distance market for the first time.[109]

After nearly a century of regulation that did more to prevent competition than to serve consumer welfare, Washington policymakers finally began taking steps to repeal some of the most egregiously anti-competitive regulations in the 1970s.

Although we should be cautious about overstating the similarity between past regulatory schemes and proposals for "network neutrality" regulation, the failures of past regulatory schemes should cause today's policymakers to be wary of enacting new regulations that

Regulated firms cultivate cozy relationships with regulators and use their influence to limit competition and raise prices.

29

could cause similar problems in the future.

Lessons of Deregulation

The most fundamental lesson of 20th-century regulatory policy is to beware of unintended consequences. It is unlikely that in any of the three cases we have just examined, Congress intended to bring about the wasteful and anti-consumer results that actually occurred. The Interstate Commerce Act produced a decade of uncertainty as the courts struggled to interpret its provisions. The ICC did little to control the railroads over the subsequent two decades. Then, from 1920 until the late 1970s, it pursued policies that overtly aided railroads, trucking companies, and other politically connected interest groups at the expense of the general public. Those outcomes surely were not what Congress had in mind in 1887.

Perhaps Congress should have known better by the time it created the Civil Aeronautics Board in 1938. Congress felt that restricting competition in the airline industry would aid the establishment of a new, struggling industry.[110] Whatever the merits of this initial rationale, the CAB continued to restrict competition long after flying had become a mature industry fully capable of standing on its own two feet. Moreover, the inability to cut prices caused airlines to engage in wasteful activities in order to attract customers. Prohibited from competing on price, airlines competed instead by spending lavishly on capacity expansion. Airlines purchased more planes and flew their routes more frequently. As a result, by 1971, fewer than half the seats on an average airline flight were filled. Rather than allowing price reductions to fill those empty seats, the CAB worked to reduce the number of flights. Airlines, now unable to compete on either price or convenience, responded by initiating the so-called "lounge wars":

> On wide-bodied aircraft, lounges were introduced in first class, then in coach. When American installed piano bars, TWA countered with electronic draw-

poker machines. Live entertainment proliferated, with musicians, magicians, wine-tasters, and Playboy bunnies.[111]

Congress hoped that the creation of the Civil Aeronautics Board would create a more rational, efficient structure for the airline industry. But the reality was just the opposite: the perverse incentives of the regulatory process caused airlines to undertake wasteful activities they never would have undertaken in a competitive market.

Bureaucratization

Another lesson of these examples is the capacity of regulators to bureaucratize an industry. The original Interstate Commerce Act included nondiscrimination language strikingly similar to today's network neutrality proposals. The problem was that applying seemingly simple rules to the real world—with hundreds of railroads, routes, and categories of cargo—was much more complicated than anyone had expected. It wasn't always clear how to apply the broad language of the ICA to specific cases, and the rapid pace of change in the railroad industry made it a challenge for the ICC to keep up.

The commission responded to these challenges in three ways. First, it pressed Congress for more power and resources. Congress repeatedly expanded the ICC's power during the early 20th century and gave it stronger powers to punish railroads that failed to comply with its orders. The ICC's staff swelled from 104 people in 1890 to 527 in 1909. By 1909, the commission was fielding 1,097 formal complaints and roughly 4,500 informal ones per year.[112]

Second, ICC regulation increasingly shifted from a complaint-based regulatory process to one based on prior approval of rate changes. During the ICC's early years, the railroads were free to set their own rates and the ICC and the courts would rule on the appropriateness of those rates after the fact. Increasingly, however, the ICC acquired the power to review rates before they would go into effect, preempting those rate changes it felt were inappropriate.

This forced the railroads to slow down to the ICC's stately pace of decisionmaking, which was the only way the commission could hope to keep up with the vast and dynamic railroad marketplace.

Finally, the ICC felt constrained by political realities not to disturb long-established pricing policies, even those that seemed to run counter to the commission's broader policy goals. Historians Ari and Olive Hoogenboom write that in order to reform the deeper problems with the railroads' rate structures, it would have needed to "disturb many powerful elements of the economy." And this it was unwilling to do.[113]

It's easy to imagine a similar fate befalling the FCC should Congress give it authority to police Internet routing policies. The Internet is made up of thousands of privately owned networks that interact with one another in a variety of complex ways. Filing a discrimination complaint with the FCC could become a common tactic in business disputes between network providers. There are enough ambiguities in the basic concept of nondiscrimination (some of which are discussed below) that it would not be difficult for a determined firm to find a plausible example of discrimination in its adversary's policies. And given the complexity of modern network infrastructure, the FCC would often be required to conduct extensive investigations to determine which complaints had merit.

It is likely that an overwhelmed FCC would respond to its overflowing docket much the same way the ICC responded to the flood of discrimination complaints a century ago: by pushing a lot of the work onto the firms they regulate. The commission would pressure ISPs to standardize their business practices and network configurations to make it easier to judge whether Commission rules were being followed. Carriers might be required to make detailed filings describing their network architecture and routing policies, and to file notices with the FCC if these changed.

Mission Creep

Once the apparatus of regulatory control has been put in place, there will be an almost irresistible temptation to use it for purposes beyond those in the original legislation or to expand its scope to new sectors of the economy. The FCC, CAB, and ICC all used their power over pricing to compel the firms they regulated to subsidize favored customers at the expense of nonfavored customers. The FCC would become embroiled in messy debates about the filtering of spam, viruses, pornography, copyrighted material, gambling applications, and other issues. The FCC's decisionmaking would likely be influenced by considerations not specifically mentioned in statute. For example, a conservative majority of the FCC might go easy on ISPs that tried to filter out content that the majority regarded as immoral, such as pornography.

There is ample precedent for this kind of political manipulation in the FCC's recent decisions. FCC chairman Kevin Martin has long been on a crusade to require "à la carte" pricing of cable television channels. As a *New York Times* columnist describes the situation, "Mr. Martin has long said that he favors à la carte because it's pro-consumer, but most people in the cable industry . . . are convinced that he favors it [because] it will allow parents to keep MTV and its ilk out of their homes."[114] The cable industry has attempted to mollify Chairman Martin by introducing a "family tier" that leaves out channels that social conservatives find objectionable.[115]

We have seen how the ICC's authority, which initially extended only to the railroad industry, was expanded to the entire surface transportation industry in the early 1930s. By the same token, once the FCC had gotten comfortable in its role as Internet neutrality cop, it might seek expanded authority to regulate the "neutrality" of search engines, operating systems, middleware platforms, e-commerce services, and the like.

Thwarted Competition and Innovation

Another clear lesson from the histories of the ICC and the FCC is that regulations can often be a powerful tool in the hands of incumbents to keep out new competitors. As discussed previously, the railroads lobbied to

Applying seemingly simple rules to the real world was much more complicated than anyone had expected.

extend the ICC's jurisdiction to the trucking industry in order to limit competition from that sector. Similarly, the FCC slowed the introduction of microwave communications technologies in order to protect AT&T's long-distance business from competition. In recent years, the cable industry has used franchising law as a weapon against Baby Bells seeking to offer competing video services, arguing that AT&T and Verizon must submit to onerous city-by-city franchising requirements before being allowed to offer video service in their existing service territories.[116]

In each case, the incumbent firm has made the plausible argument that fairness requires regulatory parity. But in practice, parity is inevitably more burdensome to the new entrant than to the incumbent. There are typically economies of scale to regulatory compliance, and experienced players generally have an advantage in practicing before a regulatory agency.

More important, complying with old regulations often constrains a new entrant's technological options. New entrants often achieve competitive advantages by rapidly deploying new, lower-cost technologies. To the extent that regulation requires all market participants to roll out "gold plated" services in an orderly fashion, it will inevitably redound to the benefits of incumbents. MCI would have rolled out its long distance service differently if it hadn't been forced to spend its first decade begging the FCC for permission to compete.

A current example of this is the regulatory challenges faced by Vonage, the pioneering Internet telephony firm. One of the ways that telephone incumbents have slowed Vonage's progress is by lobbying for strict enforcement of a wide variety of regulatory requirements that already apply to incumbent telecom firms. In recent years, the FCC, at the urging of the Baby Bells, has demanded that Vonage offer emergency 911 service,[117] redesign their networks to facilitate government eavesdropping under CALEA,[118] and pay into the Universal Service Fund that subsidizes rural telephone access.[119] Vonage has also had to beat back demands that it file paperwork with

regulators in states like New York, Texas, and Minnesota.[120]

There may be good policy arguments for each of these requirements, but the combined result has been to force Vonage executives to spend a great deal of time in federal and state courts and before federal and state regulators rather than focusing on their business. In addition, the expenses of compliance—and of filing paperwork to demonstrate compliance—have limited Vonage's ability to compete on price, which would otherwise be a key competitive advantage. Any system of regulation, no matter how well-intentioned, inevitably creates barriers to entry that hurt small challengers more than large incumbents. And limiting competition ultimately hurts consumers.

This is a particular reason for concern because there are a variety of wireless technologies on the drawing board—including WiMax and the recently completed 700 MHz wireless auction—that have the potential to shake up the market for residential broadband service. It's conceivable that some of the competitors could be small, entrepreneurial firms like MCI circa 1965. The broadband incumbents will certainly take every opportunity to place regulatory obstacles in the path of these new firms. Network neutrality rules could be turned into just such a barrier.

The advocates of network neutrality regulations mean well. But history suggests that good intentions are not sufficient to ensure that a regulatory regime will serve, rather than hinder, competition and innovation. Problems are particularly likely when, as in this case, the rules under consideration are complex and ambiguous.

The Ambiguity of Neutrality Regulation

All the disadvantages of network neutrality regulation discussed above are increased by the inherent fuzziness of the concept.[121] The Internet is sufficiently complicated and fast-changing that reasonable people disagree about exactly how to apply the concept in par-

Once the FCC got comfortable in its role as Internet neutrality cop, it might seek expanded authority to regulate the "neutrality" of search engines, operating systems, middleware platforms, e-commerce services, and the like.

ticular situations. History demonstrates that when Congress enacts a new regulatory regime, it typically leads to rulemaking and associated litigation that can drag on for the better part of a decade. In a marketplace that is evolving as rapidly as the online world, such delays can impose significant costs.

Snowe-Dorgan

The network neutrality legislation that has come closest to being approved by Congress is the Internet Freedom Preservation Act of 2006, sponsored by Sen. Olympia Snowe (R-ME), Sen. Byron Dorgan (D-ND), and others.[122] It provided that a broadband provider could not "block, interfere with, discriminate against, impair, or degrade the ability of any person to use a broadband service to access, use, send, post, receive, or offer any lawful content, application, or service made available via the Internet." It also prohibited restrictions on device attachment, special treatment for affiliated content, and charging different rates for different types of content, applications, or services. Snowe-Dorgan included exceptions for network security and parental controls.

All of these terms are rife with ambiguities. Such ambiguity is problematic when violations of network neutrality carry stiff legal penalties. It's important that those subject to the law clearly understand what the law requires of them. Applying concepts in Snowe-Dorgan to Comcast's filtering of BitTorrent illustrates this well.

BitTorrent Filtering

As discussed above, Comcast's network recently dealt with congestion by transmitting packets that misled BitTorrent and other file-sharing applications into thinking that the computer at the other end of the connection had hung up. Would the FCC have found this policy to be a violation of Snowe-Dorgan? It seems likely, but far from certain, that it would. While Comcast wasn't technically blocking any packets, Comcast's actions clearly had the effect of "interfering with" and "degrading" BitTorrent traffic.

On the other hand, BitTorrent is widely used for copyright infringement. Given that Snowe-Dorgan only protects "lawful" content, the FCC might have permitted Comcast's policy based on a showing that the vast majority of BitTorrent traffic consisted of copyrighted works. On the other hand, the FCC might have been swayed by the argument that BitTorrent users consumed a disproportionate share of traffic and that Comcast's actions were necessary to maintain the quality of other users' Internet experience.

In any event, it seems probable that the FCC's decisionmaking process would have been overtaken by events. Comcast has already announced changes to its filtering policies that would have likely rendered any ongoing proceedings moot. There is little point in having a regulatory process that moves so slowly that its decisions are irrelevant by the time they are announced.

Verizon and DNS

Another example of ISP activity that some have characterized as a network neutrality violation is Verizon's policy of redirecting failed DNS queries to Verizon's own search engine. As discussed previously, DNS servers translate a domain name (such as cato.org) into a corresponding IP address. When a user attempts to access a domain name that does not exist (perhaps because of mistyping) a DNS server is expected to return an error message and allow the application to decide how to handle the error. Instead, Verizon's DNS servers return the IP address of its own search engine, allowing Verizon to generate some ad revenue.

As we've already noted, DNS servers are just another network endpoint, architecturally speaking. Users who are dissatisfied with the behavior of Verizon's DNS server are free to use a different one. There are a number of reasons to criticize Verizon's DNS policy, but network neutrality doesn't seem to be among them.

Some experts disagree. When Ed Felten discussed the incident on his *Freedom to Tinker* blog, he characterized Verizon's actions as "a more clear-cut neutrality violation" than Comcast's interference with BitTorrent, because

History suggests that good intentions are not sufficient to ensure that a regulatory regime will serve, rather than hinder, competition and innovation.

Verizon is "interfering with the behavior of the DNS protocol."[123] Comments on the post by Felten's tech-savvy readers were evenly divided on whether Verizon's actions implicated network neutrality.

How would the FCC have applied Snowe-Dorgan in this example? The language of the bill doesn't do much to answer the question. The only thing we can predict for sure is that a lot of lawyers would have been involved.

Free WiFi

It is also unclear who would be subject to network neutrality rules. Snowe-Dorgan defined a "broadband service provider" as any person who "controls, operates, or resells and controls any facility used to provide broadband service to the public, whether for a fee or for free." This seems to suggest that coffee shops, hotels, and other businesses that offered WiFi access as an incidental part of their business would be subject to network neutrality requirements. If such a provider happened to run a poorly configured firewall, for example, such a business could be hauled before the FCC to justify its network configurations.

There is no good policy rationale for subjecting every small consumer business in America to network neutrality rules. And of course, the FCC is unlikely to go out of its way to harass small businesses. But the literal meaning of Snowe-Dorgan would have it do so. Had it passed, the FCC would inevitably have received a complaint about a small business's WiFi service, and the commission would have had to issue rules about who is subject to network neutrality regulations.

IPTV

Snowe-Dorgan explicitly exempts any service regulated under Title VI of the Communications Act from network neutrality regulations. This is the section governing cable television. That suggests a potential loophole for network owners wishing to skirt network neutrality rules: offer video as part of the service and characterize it as a "cable service" rather than a broadband service. Cable and telephone companies might use this cable

loophole to do many of the things that concern network neutrality proponents. Digital cable services already have video-on-demand services, digital channel guides, and picture-in-picture support. They could syndicate content from an Internet-based video service like YouTube, or roll out enhanced digital services such as video games, to evade the spirit of legislation like Snowe-Dorgan.

That would put the FCC in the awkward position of deciding how much functionality a cable system can have before it becomes a full blown broadband service. Snowe-Dorgan relies on the definition of cable television found in the 1996 Telecommunications Act. Given the rapid pace of technological progress since then, there are good reasons to doubt if this definition would be up to the task, and if not, another years-long inquiry by the FCC would have to be undertaken.

Jitter

As previously discussed, random delays in packet delivery (called "jitter") degrade the performance of latency-sensitive applications. Of course, some of the major broadband providers are also telephone companies, and these firms may be tempted to increase the jitter of their networks in order to discourage competition from VoIP services. Such a strategy would sidestep some of the difficulties that would come with a strategy of explicit packet filtering because it could be applied indiscriminately to all traffic without significantly degrading the quality of non–latency-sensitive applications such as websites and e-mail. On the other hand, it *would* degrade the quality of latency-sensitive applications like network gaming and remote terminal sessions, so the strategy would not be without collateral damage.

In either event, Ed Felten has pointed out, this could be an especially difficult case for regulators to deal with.[124] Some networks have jitter for reasons beyond the control of the network owner. In other cases, jitter may have innocent explanations, but network owners may choose not to perform network upgrades that would reduce it. In still other cases, a network owner might deliberately

Ambiguity is problematic when violations of network neutrality regulations carry stiff legal penalties.

introduce jitter but pretend it had made the change that caused it for unrelated reasons.

It could be quite difficult for a regulator to distinguish among these cases. Of course, a network owner under a network neutrality regime will never admit that it is increasing jitter on its network. So the FCC could be forced to second-guess the complex network-management decisions of network owners.

The Consequences of Ambiguity

It took 10 years for the Supreme Court finally to resolve questions about the Interstate Commerce Commission's authority in 1897.[125] MCI had to wait close to a decade for permission to build a competitive long-distance network. Things haven't gotten any better in recent years. As Christopher Yoo has pointed out,[126] the 1996 Telecommunications Act prompted a flurry of legal wrangling before the FCC and the federal courts, which culminated in the Supreme Court's 2002 decision *Verizon Communications Inc. v. FCC*[127] and its 2005 decision in *National Cable and Telecommunication Association v. Brand X Internet Services*.[128]

During this nine-year period, neither incumbent firms nor potential challengers knew what rules would govern any new infrastructure investments they might make. This made incumbents less likely to upgrade their facilities. But it was much worse for competitive firms whose business plans depended on the outcome of these cases. Because many such firms were bleeding red ink, the incumbents needed only to drag out the proceedings long enough for the new entrants to run out of money.

The same could easily happen if Congress enacted network neutrality regulations. The FCC would likely receive a flood of complaints about the behavior of various network owners. It would take months, if not years, for the FCC to rule on these complaints, and many of them would then be appealed to the courts. At best, this would be a distraction for firms that ought to be focusing on developing innovative new products. At worst, the lack of clarity could cause some firms to delay entry into the market until the uncertainty had been resolved.

Economic efficiency requires clear legal rules. When rules are overly complex or ambiguous, entrepreneurs are forced to spend time on unproductive activities like lobbying and litigation, instead of on serving their customers. The Snowe-Dorgan bill was full of unnecessary ambiguity and complexity that would have forced high-tech companies to hire lobbyists and lawyers instead of engineers. Given the lessons of history, we are fortunate that Congress did not enact such regulations into law.

Conclusion

For all the passionate disagreement that has characterized the network neutrality debate in recent years, there may be fewer differences than either side is willing to admit. Both sides hail the rapid growth of the Internet and the fiercely competitive online marketplace it has produced. Both seek to prevent a return to the monopolistic communications market of past decades, in which large companies and government regulators colluded to maintain the status quo to the detriment of consumers.

Yet many deregulationists underestimate the importance of the Internet's end-to-end architecture and are too cavalier about abandoning the neutral network for a tiered, filtered, more centrally managed one. The decentralization made possible by the Internet's open architecture is the key to its astonishing growth, and there is little reason to think that it would be improvement for the Internet's decentralized "dumb" architecture to be replaced by a more centralized "smart" one.

For their part, the "openists" are unduly pessimistic about the durability of the open networking architecture they have championed for the last quarter century. In the 1980s, the Internet triumphed over proprietary networks precisely because the partisans for open networks were right about the fundamental advantages of open technologies. Now that the Internet is the world's dominant communications network, those same fundamental advan-

The only thing we can predict for sure is that a lot of lawyers would be involved.

> **Any effort to introduce centralized control over the Internet will be stymied by the simple fact that centralized control is inefficient.**

tages will make end-to-end extremely difficult to dislodge. Any effort to introduce centralized control over the Internet will be stymied by the simple fact that centralized control is inefficient. That, along with the vigilance of rank-and-file partisans for open networks, will be sufficient to maintain the Internet's open architecture.

Only one institution in American society has the size and power to bring about a return to the bad old days of monopolistic communications markets: the federal government. Government regulation of private industry frequently leads to unintended consequences, and industry incumbents often find ways to turn the regulatory system to their own benefit. It would be unfortunate if a hasty effort to enact network neutrality rules led to decades of litigation and regulatory battles over the meanings of network neutrality concepts when the focus should be on developing new and better technology. And it would be especially ironic if, in their effort to protect the Internet against centralized control by major telecom companies, the openists laid the groundwork for a regulatory regime that telecom incumbents ultimately used to limit competition in the broadband industry.

Notes

1. Tim Wu, "The Broadband Debate: A User's Guide," *Journal of Telecommunications and High Technology Law* 3, no. 1 (2004).

2. Tim Berners-Lee is no relation to the author.

3. See, for example, Alfred E. Kahn, "Network Neutrality" (Working Paper no. RP07-05, AEI-Brookings Joint Center, March 2007), http://ssrn.com/abstract=973513.

4. Paul Kouroupas, vice president for regulatory affairs and security officer at Global Crossing, a "tier 1" Internet backbone operator, submitted comments to the FCC in response to Vuze Inc.'s "Petition to Establish Rules Governing Network Management Practices by Broadband Network Operators" (Docket no. WC 07-52). Kouroupas noted that the Internet backbone has traditionally been beyond the jurisdiction of the FCC and that this flexibility will be needed as the backbone providers rapidly upgrade their facilities to cope with a flood of high-definition video traffic.

5. A protocol is a language that two devices use to communicate over a computer network.

6. Katie Hafner and Matthew Lyon, *Where Wizards Stay Up Late: The Origins of the Internet* (New York: Simon and Schuster, 1996), pp. 192–93.

7. Ibid., p. 194.

8. TCP and IP stand for Transmission Control Protocol and Internet Protocol, respectively, but the phrase "TCP/IP protocols" refers to a broader suite of networking protocols, including UDP (User Datagram Protocol), ICMP (Internet Control Message Protocol), and others, which form the foundation of the Internet.

9. The initial version of the protocol that emerged from a seminar at Stanford held in 1973 proposed a single protocol called TCP that would be responsible for ensuring that packets were transmitted reliably from source to destination. A subsequent revision, developed in 1978, split the original protocol into two layers: an IP layer that was responsible for delivering individual packets, and a TCP protocol that was responsible for error correction. In addition to simplifying the process of implementing TCP/IP on a variety of different networks, this had the added benefit of allowing applications that didn't need the error-correction features of the TCP layer to avoid the overhead by using lighter-weight protocols such as UDP. See Janet Abbate, *Inventing the Internet* (Cambridge, MA: MIT Press, 1999), pp. 127–30.

10. J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-End Arguments in System Design," *Second International Conference on Distributed Computing Systems*, Paris, France, April 8–10, 1981, pp. 509–512.

11. At its inception in late 1990, the project team consisted of Berners-Lee, his colleague Robert Cailliau, and a student, Nicola Pellow, with part-time assistance from Bernd Pollermann. Berners-Lee and Cailliau "argued tirelessly for resources from CERN but never got quite what they asked for." As a result, the team remained woefully undermanned, and much of the Web's early development was due to volunteer contributions from outside of CERN, most notably NCSA's Mosaic browser. See James Gillies and Robert Cailliau, *How the Web Was Born* (New York: Oxford University Press, 2000), pp. 199–235.

12. Tim Wu, "Wireless Carterfone," *International Journal of Communication* [online] 1 (2007): 389.

13. David Pogue, "The iPhone Matches Most of Its Hype," *New York Times*, June 27, 2007, http://www.nytimes.com/2007/06/27/technology/circuits/27pogue.html.

14. Miguel Helft and John Markoff, "Google Enters the Wireless World," *New York Times*, No-vember 5, 2007, http://www.nytimes.com/2007/11/05/technology/05cnd-gphone.html.

15. Laura M. Holson, "Verizon Plans Wider Options for Cellphone Users," *New York Times*, November 28, 2007, http://www.nytimes.com/2007/11/28/technology/28phone.html.

16. Robert Hahn and Scott Wallsten, "The Eco-nomics of Net Neutrality." *Economists' Voice* 3, no. 6, (June 2006), http://www.aei-brookings.org/publications/abstract.php?pid=1067.

17. An excellent overview of the history and eco-nomics of interconnection on the Internet is Geoff Huston, "Interconnection, Peering, and Settle-ments," (address, The Internet Society's, INET'99 conference, San Jose, CA), http://www.isoc.org/inet99/proceedings/1e/1e_1.htm.

18. Christopher S. Yoo, "Network Neutrality and the Economics of Congestion," *Georgetown Law Journal* 94, no. 6 (August 2006): 1873.

19. See, for example, Gigi Sohn, "Time Warner Steps Up to the Plate on Bandwidth Usage," post-ing on Public Knowledge's Policy Blog, January 17, 2008. http://www.publicknowledge.org/node/1356, and Adam Thierer, "Broadband Metering Experiment in the Works in Texas?" Technology Liberation Front, January 16, 2008, http://techliberation.com/2008/01/16/broadband-metering-experiment-in-the-works-in-texas/.

20. Christopher S. Yoo, "Beyond Network Neutral-ity," *Harvard Journal of Law and Technology* 19, no.1 (Fall 2005): 21–22.

21. Edward W. Felten, "Verizon Violates Net Neu-trality with DNS Deviations," *Freedom to Tinker*, November 12, 2007, http://www.freedom-to-tinker.com/?p=1227.

22. Benjamin Teitelbaum and Stanislav Shalunov, "Why Premium IP Service Has Not Deployed (and Probably Never Will)," Internet2 QoS Working Group, Informational Document, May 3, 2002, http://qbone.internet2.edu/papers/non-architectural-problems.txt.

23. S. Blake et al, "An Architecture for Differenti-ated Services," Request for Comments: 2475, December 1998, Network Working Group, Inter-net Engineering Task Force, http://www.ietf.org/rfc/rfc2475.txt.

24. A good overview of this architecture can be found in "Diffserv—The Scalable End-to-End Quality of Service Model," Technology White Paper, Cisco Systems, August 2005, http://www.ieng.net/en/US/technologies/tk543/tk766/technologies_white_paper09186a00800a3e2f.html.

25. J. L. Adams, L. G. Roberts, and A. Ijsselmuiden, "Changing the Internet to Support Real-time Content Supply from a Large Fraction of Broad-band Residential Users," *BT Technology Journal* 23, no. 2 (April 2005).

26. Yoo, "Network Neutrality and the Economics of Congestion," p. 1907.

27. Ibid., p. 1854.

28. While the point is tangential to the paper, I can't resist pointing readers to John Gilmore's provocative but surprisingly persuasive argument against server-side spam filtering. Gilmore's argu-ment goes beyond the position taken here, con-tending that the end-to-end principle should be applied to the routing policy of mail servers as well as Internet routers. This is, to be clear, a broader conception of network neutrality than I advocate in this paper. See "Verio Censored John Gilmore's Email under Pressure from Anti-spam-mers," http://www.toad.com/gnu/verio-censorship.html.

29. *Internet Freedom Preservation Act*, S. 2917, 109th Cong., 2nd Sess., http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=109_cong_bills&docid=f:s2917is.txt.pdf.

30. Robert W. Hahn and Robert E. Litan, "The Myth of Network Neutrality and What We Should Do About It," AEI-Brookings Joint Center for Regulatory Studies, related publication 06-33, No-vember 2006, http://ijoc.org/ojs/index.php/ijoc/article/view/161/87.

31. Robert Braden, David Clark, and Scott Shenker, "Integrated Services in the Internet Architecture: An Overview," RFC 1633, June 1994, http://www.ietf.org/rfc/rfc1633.txt

32. RFC 1633 has been supplanted by RFC 2205 and RFC 3209, and these protocols are sometimes used for private IP-based virtual networks, but RFC 1633 and its successors have not been wide-ly deployed on the public Internet.

33. Vinton Cerf, Yogen Dalai, and Carl Sunshine, "Specification of Internet Transmission Control Program." RFC 675, December 1974, http://www.ietf.org/rfc/rfc675.txt.

34. Information Sciences Institute, "Internet Proto-col," RFC 791, September 1981, http://www.ietf.org/rfc/rfc791.txt.

35. Vinton Cerf, "Pre-emption," RFC 794, Septem-ber 1981, http://www.ietf.org/rfc/rfc794.txt.

36. I am indebted to the blogger "Cog" (pseud.), whose blog post is the basis for this paragraph. See "Notes on (Hahn Litan 06): Network Neutrality Part 1: Requests For Comments," The Abstract Factory, January 24, 2007, http://abstractfactory .blogspot.com/2007/01/notes-on-hahn-litan-06-network.html.

37. Bill D. Herman, for example, writes that "The broadband providers' candor regarding their intention to begin discriminating should be proof enough that today's generally discriminatory Internet is in danger." He does not seriously entertain the possibility that network owners might lack the power to undermine the Internet's end-to-end architecture. See "Opening Bottlenecks: On Behalf of Mandated Network Neutrality." *Federal Communications Law Journal* 59, no. 1 (2006): 127. Paul Misener, vice president of Global Public Public Policy at Amazon.com, wrote in a December 2, 2002, FCC filing that it was "highly likely" that broadband service providers would impair delivery of content based on "easily obtainable knowledge of the source and nature" of that content. See Amazon.com filing in *Matter of Appropriate Regulatory Treatment for Broadband Access to the Internet over Cable Facilities*, FCC CS Docket no. 02-52; FCC 02-77.

38. David Chartier, "Preliminary iPhone 1.1.1 jailbreak announced," Ars Technica, October 8, 2007, http://arstechnica.com/journals/apple.ars/2007/10/08/preliminary-iphone-1-1-1-jailbreak-announced.

39. Edward W. Felten, "AACS Plays Whack-a-Mole with Extracted Key," Freedom to Tinker, May 1, 2007, http://www.freedom-to-tinker.com/?p=1152.

40. Jay Adelson, "What's Happening with the HD-DVD Stories?" Digg the Blog, May 1, 2007, http://blog.digg.com/?p=73.

41. Kevin Rose, "Digg This: 09-f9-11-02-9d-74-e3-5b-d8-41-56-c5-63-56-88-c0," Digg the Blog, May 1, 2007, http://blog.digg.com/?p=74.

42. Matt Hines, "Yahoo IM Update Shuts out Third Parties," CNet News.com, September 17, 2003, http://www.news.com/2100-1012-5078361.html; Jim Hu, "Yahoo Walls out Trillian," CNet News.com, September 26, 2003, http://www.news.com/Yahoo-walls-out-Trillian/2100-1032_3-5082812.html; Jim Hu, "Yahoo to Trillian: Talk to the Hand," CNet News.com, June 23, 2004, http://www.news.com/2100-1032_3-5245821.html.

43. Lisa M. Bowman, "AOL Blocks Instant Messaging Start-up," CNet News.com, January 30, 2002, http://www.news.com/2100-1023-8266 25.html; Jim Hu, "MSN Messenger Upgrade Blocks Trillian," CNet News.com, August 20, 2003, http://www.news.com/2100-1032_3-50664 12.html.

44. Clay Shirky, *Here Comes Everybody: The Power of Organizing without Organizations* (New York: Penguin Press: 2008).

45. Roger O. Crockett, "At SBC, It's All About 'Scale and Scope,'" *Business Week* online extra, November 7, 2005, http://www.businessweek.com/@@n34h *IUQu7KtOwgA/magazine/content/05_45/b395 8092.htm.

46. Abbate, p. 84.

47. Ibid., p. 85.

48. Ibid., pp. 136–7.

49. Tim Wu, "Network Neutrality, Broadband Discrimination," *Journal of Telecommunications and High Technology Law* 2 (2003).

50. For example, dreamhost.com advertises web and e-mail hosting, shell access, and 500 GB of storage space for $9.95 per month with a one-year contract.

51. Jacqui Cheng, "Evidence Mounts that Comcast is Targeting BitTorrent Traffic," Ars Technica, October 19, 2007, http://arstechnica.com/news.ars/post/20071019-evidence-mounts-that-comcast-is-targeting-bittorrent-traffic.html.

52. This was reported by user "funchords" (pseud.) on the DSLReports forums on May 12, 2007. See http://www.dslreports.com/forum/r18323368-Comcast-is-using-Sandvine-to-manage-P2P-Connections.

53. This was reported by user "Deluxe05" (pseud.) on the DSLReports Forums on November 11, 2007. See http://www.dslreports.com/forum/r19 386386-Comcast-Sandvine-and-the-latest-WoW-patch-v230.

54. Eric Bangeman, "Comcast Traffic Blocking: Even More Apps, Groupware Clients Affected," Ars Technica, October 21, 2007, http://arstechnica.com/news.ars/post/20071021-comcast-traffic-blocking-even-more-apps-groupware-clients-affected.html.

55. Comcast Website, "FAQ," retrieved on February 12, 2008 from http://www.comcast.net/help/faq/index.jsp?faq=Hot118988.

56. Anne Broache, "Verizon: No 'Need' to Degrade P2P Traffic . . . Yet," CNet News.com, February 11, 2008, http://www.news.com/8301-10784_3-986 9327-7.html.

57. "Ernesto" (pseud.), "Encrypting BitTorrent to Take out Traffic Shapers," TorrentFreak, February 5, 2006, http://torrentfreak.com/encrypting-bittorrent-to-take-out-traffic-shapers/.

58. Ibid., "How to Bypass Comcast's BitTorrent Throttling," TorrentFreak, October 21, 2007, http://torrentfreak.com/how-to-bypass-comcast-bittorrent-throttling-071021/.

59. Brad Stone, "Comcast Adjusts Way It Manages Internet Traffic," *New York Times*, March 28, 2008, http://www.nytimes.com/2008/03/28/technology/28comcast.html.

60. The FCC's decision, coming as this paper was going to press, raises important legal issues that are beyond its scope.

61. Joel Hruska, "DOCSIS 3.0, Possible 100Mbps Speeds Coming to some Comcast Users in 2008," *Ars Technica*, November 30, 2007, http://arstechnica.com/news.ars/post/20071130-docsis-3-0-possible-100mbps-speeds-coming-to-some-comcast-users-in-2008.html.

62. Crockett.

63. Yochai Benkler, *The Wealth of Networks: How Social Production Transforms Markets and Freedom* (New Haven: Yale University Press, 2006).

64. Ibid., p. 158.

65. Ibid., pp. 155–6.

66. Ibid., p. 156.

67. Ibid.

68. http://www.techliberation.com/.

69. Because peering agreements are confidential, it is difficult to assemble a definitive list of Tier 1 network providers, but as of early 2008, Wikipedia listed the following as Tier 1 carriers: AOL Transit Data Network, AT&T, Global Crossing, Level 3, Verizon, NTT, Qwest, SAVVIS, and Sprint. See http://en.wikipedia.org/wiki/Tier_1_carrier. An excellent introduction to peering in *Ars Technica* estimated that there are seven Tier 1 providers. Rudolph van der Berg, "How the 'Net Works: An Introduction to Peering and Transit," *Ars Technica*, September 2, 2008, http://arstechnica.com/guides/other/peering-and-transit.ars/.

70. Benkler, p. 158.

71. Ibid. p. 240.

72. The Associated Press confirmed reports that Comcast was interfering with BitTorrent traffic using such tests. Peter Svensson, "Comcast Blocks Some Internet Traffic," *Washington Post*, October 19, 2007, http://www.washingtonpost.com/wp-dyn/content/article/2007/10/19/AR2007101900842.html.

73. Mike Masnick, "Since When Is It Illegal to Just Mention a Trademark Online?" Techdirt, January 5, 2005, http://www.techdirt.com/articles/20050105/0132239.shtml.

74. Markos Moulitsas, "Liberal Blogger Blocked from Kentucky State-Owned Computers," Daily Kos, June 21, 2006, http://www.dailykos.com/story/2006/6/21/111252/975.

75. Adam Liptak, "Verizon Reverses Itself on Abortion Messages," *New York Times*, September 27, 2007, http://www.nytimes.com/2007/09/27/business/27cnd-verizon.html.

76. Lawrence Lessig made this point in a presentation at Stanford University on April 17, 2008. His talk can be heard at http://lessig.org/blog/2008/04/testifying_fcc_stanford.html.

77. For example, Michelle Kessler, "Internet Fast Lane Plan Worries Small Companies," *USA Today*, June 18, 2006.

78. Matthew Broersma, "AOL 4.0 Assimilates the Net," ZD Net, Sep 26, 1998, http://news.zdnet.com/2100-9595_22-512081.html.

79. Elinor Mills, "AOL: You've Got Free E-mail." CNet News.com, August 2, 2006, http://www.news.com/AOL-business-model-in-transition/2100-1025_3-6101144.html.

80. Yoo, "Network Neutrality and the Economics of Congestion," p. 1851.

81. ESPN, generally cited as the most expensive cable channel, cost more than $2.50 per subscriber in 2004. See Sean Gregory, "Why ESPN Is The Crown Jewel," *Time*, February 23, 2004.

82. Internet Freedom Preservation Act, S. 2917, 2006.

83. For an overview of the progressive legal theories that made increased government regulation possible, see Richard Epstein, *How Progressives Rewrote the Constitution* (Washington: Cato Institute, 2006), pp. 52–110.

84. Gabriel Kolko, *Railroads and Regulation, 1877–1916* (Princeton, NJ: Princeton University Press), 1965, pp. 47–48.

85. Ari and Olive Hoogenboom, *A History of the ICC: From Panacea to Palliative* (New York: W.W. Norton, 1976), p. 26.

86. Ibid., p. 30.

87. *Interstate Commerce Commission v. The Cincinnati, New Orleans and Texas Pacific Railway Company et al.*, 167 U.S. 479 (1897).

88. Hoogenboom, p. 39.

89. Ibid., pp. 44, 52, 60.

90. Ibid., pp. 55–6.

91. Ibid., pp. 94–6.

92. Ibid., pp. 130–1.

93. Ibid., p. 137.

94. Robert C. Fellmeth, *The Interstate Commerce Omission, the Public Interest and the ICC: The Ralph Nader Study Group Report on the Interstate Commerce Commission* (New York: Viking Press, 1970).

95. "Civil Aeronautics Board Policy: An Evaluation," *Yale Law Journal* 57, no. 6 (April 1948).

96. For example, Adam Thierer has argued that federal and state regulation was essential to the consolidation of the Bell monopoly. See Adam D. Thierer, "Unnatural Monopoly: Critical Moments in the Development of the Bell System Monopoly," *Cato Journal* 14, no. 2 (Fall 1994).

97. John Brooks, *Telephone: The First Hundred Years* (New York: Harper and Row, 1975), p. 196.

98. Peter Temin, *The Fall of the Bell System: A Study in Prices and Politics* (Cambridge: Cambridge University Press, 1987), p. 22.

99. Ibid., p. 26.

100. Ibid., p. 29.

101. Ibid., pp. 47–54.

102. A good overview of this literature is Sam Peltzman, "The Economic Theory of Regulation after a Decade of Deregulation," *Brookings Papers on Economic Activity, Microeconomics* 1989.

103. G. Stigler, "The Theory of Economic Regulation," *The Bell Journal of Economics and Management Science* 2, no. 1 (Spring, 1971).

104. Further discussion of regulatory capture and network neutrality can be found in Adam D. Thierer, "'Network Neutrality': Digital Discrimination or Regulatory Gamesmanship in Cyberspace?" Cato Institute Policy Analysis no. 507, January 12, 2004, pp. 17–19.

105. Martha Derthick and Paul J. Quirk, *The Politics of Deregulation* (Washington: Brookings Institution, 1985), pp. 40–44.

106. Thomas Gale Moore, "U.S. Airline Deregulation: Its Effects on Passengers, Capital, and Labor," *Journal of Law and Economics* 29, no. 1 (April 1986): 2–3.

107. Ibid., p. 8.

108. John C. Taylor, "Regulation of Trucking by the States." *Regulation* 17, no. 2 (Spring 1994).

109. Temin, pp. 113–31, 217–276.

110. "Civil Aeronautics Board Policy: An Evaluation."

111. Richard H. K. Vietor, "Contrived Competition: Airline Regulation and Deregulation, 1925–1988," *Business History Review* 64, no. 1 (Spring 1990).

112. Hoogenboom, p. 55.

113. Ibid., p. 55.

114. Joe Nocera, "Bland Menu If Cable Goes à la Carte," *New York Times*, November 24, 2007.

115. Ken Fisher, "First Peek at a 'Family Cable Tier,'" *Ars Technica*, December 15, 2005, http://arstechni ca.com/news.ars/post/20051215-57 82.html.

116. Jerry Brito and Jerry Ellig, "Video Killed the Franchise Star: The Consumer Cost of Cable Franchising and Policy Alternatives," March 2006. Available at the Social Science Research Network, http://ssrn.com/abstract=893606.

117. Marguerite Reardon, "Net Phone Operators Reach E911 Deadline," CNet News.com, November 28, 2005, http://www.news.com/Net-phone-operators-reach-E911-deadline/2100-7352_3-5974196.html.

118. Declan McCullagh and Ben Charny, "Feds Back Wiretap Rules for Internet," CNet News.com, August 4, 2004, http://www.news.com/Feds-back-wiretap-rules-for-Internet/2100-7352_3-5296 417.html.

119. Anne Broache, "Appeals Court Ruling Upholds Net Phone Taxes," CNet News.com, June 1, 2007, http://www.news.com/Appeals-court-ruling-upholds-Net-phone-taxes/2100-7352_3-61882 23.html.

120. Ben Charny, "Vonage Beats Back New York Ruling," CNet News.com, June 30, 2004, http://www.news.com/Vonage-beats-back-New-York-ruling/2100-7352_3-5253841.html. Ben Charny, "States Gang Up on Vonage," CNet News.com, April 11,

2005, http://www.news.com/States-gang-up-on-Vonage/2100-1036_3-5662937.html.

121. Adam D. Thierer, pp. 6–7.

122. Internet Freedom Preservation Act, S. 2917, 2006.

123. Edward W. Felten, "Verizon Violates Net Neutrality with DNS Deviations," Freedom to Tinker, November 12, 2007, http://www.freedom-to-tinker.com/?p=1227.

124. Edward W. Felten, "Nuts and Bolts of Network Neutrality," Center for Information Technology Policy, Princeton University, July 6, 2006, http://itpolicy.princeton.edu/pub/neutrality.pdf.

125. *Interstate Commerce Commission v. The Cincinnati, New Orleans and Texas Pacific Railway Company et al.*, 167 U.S. 479 (1897).

126. Christopher Yoo, "Beyond Network Neutrality," pp. 41–42.

127. 535 U.S. 467 (2002).

128. 545 U.S. 967 (2005).

STUDIES IN THE POLICY ANALYSIS SERIES