International Journal of Communication 1 (2007), 595-606 1932-8036/20070595

net neutrality issue

The Myth of Network Neutrality and What We Should Do About It

ROBERT W. HAHN
and
ROBERT E. LITAN

AEI-Brookings Joint Center for Regulatory Studies

Executive Summary

A quarter century ago, there was a very influential paper that shaped thinking on how best to design what we now call the Internet. The article offered a design principle called "end-to-end." The idea was to keep the inner part of a computer network as simple as possible and allow the "intelligence" to reside at the edges of the network closer to the end-user.

Proponents of this grand design have pushed for net neutrality legislation, which would discourage access providers from placing any intelligence in the inner part of the network. Their ideal of a "dumb network" would be achieved by preventing access providers from charging content providers for prioritized delivery and other quality enhancements made possible by placing intelligence at the center of the network.

This essay examines the merits of the end-to-end argument as it relates to the net neutrality debate. First, we review the evidence on the current status of the Internet, concluding that all bits of information are not treated equally from an economic standpoint. Second, we demonstrate that because consumers and business place a premium on speed and reliability for certain kinds of Internet services, network owners and specialized service providers have responded with customized offerings. Third, we consider our findings in the context of the current legislative proposals involving net neutrality. Fourth, we consider some of the problems with regulating prices and quality of service, which is essentially what the net neutrality proponents propose. Our principle conclusions are that the end-to-end principle does not make sense from an economic perspective and that further regulation of the Internet is not warranted at this point in time.

Robert W. Hahn: **rhahn@communityprep.org**
Robert E. Litan: **rlitan@kauffman.org**
Date submitted: 2007-06-07

The Myth of Network Neutrality and What We Should Do About It

ROBERT W. HAHN AND ROBERT E. LITAN†

A quarter century ago, three MIT computer scientists wrote a paper that shaped the thinking on how best to design what we now call the Internet, extolling a design principle dubbed "end-to-end."[1] The idea was to keep the inner part of a computer network as simple as possible and allow the "intelligence"—that is, the ability to prioritize one data packet over another—to reside at the edges of the network closer to the end-user.[2] Several leading Internet thinkers, such as Larry Lessig of the Stanford University Law School, have suggested that end-to-end design has given rise to a huge amount of innovation in the ways the Internet is used.[3] By treating all bits of data equally, entrepreneurs at the edges of the network can compete—or so the argument goes—to bring consumers new products and applications.

Many proponents of this grand design have pushed for "net neutrality" legislation, which would discourage Internet access providers from placing any intelligence in the inner part of the network. Their ideal of a "dumb network" would be achieved simply by preventing access providers from charging content providers for priority delivery or other quality enhancements, such as guaranteed minimum bandwidth. The net neutrality debate is especially important today. Access providers are in the midst of a multibillion dollar campaign to upgrade their networks, using a mixture of capacity expansion and electronic enhancements to carry broadband content, including multiple HDTV signals. Under net neutrality regulation, they would be forced to meet this growing demand with increases in capacity only, which is a very costly solution for both access providers and their subscribers. And the prospects for legislation have improved drastically because key Democratic supporters are back in leadership positions.

Here, we review the evidence on the current status of the Internet, concluding that all bits of information are not treated equally from an economic standpoint. Second, we show that because

---

† Mr. Hahn and Mr. Litan co-founded and direct the AEI-Brookings Joint Center. The authors wish to thank Caroline Cecot, Nicholas Economides, Hal Singer, Scott Wallsten, and Molly Wells for helpful suggestions. An earlier version of this essay appeared in *The Milken Institute Review*. The views expressed here are their own.

1. *See* Jerome H. Saltzer, David P. Reed & David D. Clark, *End-to-end arguments in system design*, Second International Conference on Distributed Computing Systems 509-12 (1981); Jerome H. Saltzer, David P. Reed & David D. Clark, *End-to-end arguments in system design*, 2(4) ACM TRANSACTIONS ON COMPUTER SYSTEMS 277-88 (1984).

2. For an explanation of the intelligence capabilities of next-generation networks, see HSBC GLOBAL RESEARCH, NET NEUTRALITY: TELECOMS MUST MONETISE THE NET RATHER THAN BE TRAPPED IN IT—WE SET OUT OUR 'ABC' PATH TO FREEDOM (2006).

3. *See*, *e.g.*, *Hearing on Net Neutrality Before the S. Comm. on Commerce, Science, and Transportation*, 109th Cong. (2006) (statement of Lawrence Lessig). AEI-Brookings Joint Center Testimony 06-01. March 2006, *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1254.

consumers and businesses place a premium on speed and reliability for certain Internet services, network owners and specialized service providers have responded to market incentives with customized offerings. We think this is a good thing—not the initial step toward crony Internet capitalism, as many net neutrality advocates would argue. Third, we consider Internet regulation proposals, concluding that they are misguided and even extreme by the standards of those advocating net neutrality. In particular, the proposals would require access providers that offered enhanced quality to one content provider to make the same level of service quality available to all content providers at no extra charge. Fourth, we consider some of the problems with regulating prices and quality of service, which is essentially what the net neutrality proponents propose. Pricing flexibility is a good thing in general, and there is no reason to believe that it is not a good thing here. In our view, the antitrust laws are adequate to police access providers who have incentives to discriminate among content providers in the provision of service quality.[4]

---

[4]. *See,* for example, Robert W. Hahn & Scott Wallsten, *The Economics of Net Neutrality*, 3(6) THE ECONOMISTS' VOICE (2006), *available at*: http://www.bepress.com/ev/vol3/iss6/art8. For other economic perspectives on net neutrality and related issues, see Alfred E. Kahn, *Network Neutrality*, AEI-Brookings Joint Center Related Publication 07-05, Mar. 2007, *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1374; Benjamin E. Hermalin and Michael Katz, *The Economics of Product-Line Restrictions With an Application to the Network Neutrality Debate*, AEI-Brookings Joint Center Working Paper 07-02, Feb. 2007 *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1362; John Mayo, *Net Neutrality: The Prequel*, AEI-Brookings Joint Center Policy Matters 07-12, Mar. 2007 *available at* http://www.aei-brookings.org/policy/page.php?id=283; Bruce M. Owen, *The Net Neutrality Debate: Twenty Five Years after United States v. AT&T and 120 Years after the Act to Regulate Commerce*, AEI-Brookings Joint Center Working Paper 07-03, Feb. 2007 *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1363; Farrell, Joseph and Philip Weiser (2003), "Modularity, Vertical Integration and Open Access Policies: Towards a Convergence of Antitrust and Regulation in the Internet Age," *Harvard Journal of Law and Technology* 17(1):85-135. http://jolt.law.harvard.edu; and Bruce M. Owen and Gregory Rosston (2003), "Local Broadband Access: Primum Non Nocere or Primum Processi? A Property Rights Approach," AEI-Brookings Joint Center Related Publication 03-19, Aug. 2003, *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=285.

The Myth That All Bits Are Treated Equal

The Internet is literally a network of computers. The network moves data to and from those computers. The network includes a set of routers—think of highway interchanges—connected by long wires. Packets of data get passed from one router to another, until they arrive at their destination.[5] The packets travel at different speeds for different kinds of users, with the speed determined by a number of factors, including congestion on the network and bandwidth capacity at the point of connection. Many users subscribe to services that allow them to use the Internet whenever they want at specified maximum speeds. Thus, a typical consumer may pay $15 per month for a maximum speed of 786,000 bits per second and $30 for a maximum four times greater.[6]

The net neutrality debate is largely about whether access providers have the right to give some data preferential treatment over other data—and then charge content providers for those preferences and for other enhanced services made possible by the technological ability to discriminate among data packets. Content providers could either pass a share of that surcharge for priority delivery onto their customers or use the enhanced capacity to generate revenue in myriad other ways—say, through advertising.

The original proponents of the end-to-end principle argued that most features in the middle of a communications system are redundant if the end-user must implement them a second time.[7] This viewpoint leads to the model of a dumb network[8] with intelligence built only into its edges. According to Edward Felten, a Princeton University computer science professor who is a proponent of the end-to-end principle, the routers in the middle of the Internet "forward packets with only minor processing—all the heavy lifting takes place on the transmitting and receiving computers." [9] Felten argues that keeping things this way makes sense in large part because computers at the edges of the network are owned and controlled directly by end-users.[10]

---

[5]. For a useful technical description, see Edward W. Felten, *Nuts and Bolts of Network Neutrality*, Related Publication 06-23, Aug. 2006, *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1319.

[6]. *See* Verizon DSL Packages and Prices, *available at* http://www22.verizon.com/ForHomeDSL/channels/dsl/packages/default.asp.

[7]. *See* Jerome H. Saltzer, David P. Reed & David D. Clark, *End-to-end arguments in system design*, Second International Conference on Distributed Computing Systems 509-12 (1981). One example of a redundant feature in the middle of the network is the checksum of a file transfer, which adds up the basic components of a message and stores the resulting value.

[8]. Wikipedia credits the phrase "dumb network" to David Isenberg, a former AT&T Bell Labs employee. *See* David Isenberg, The Rise of the Stupid Network (HTML), 2.1 ACM NETWORKER (1998), *available at* http://www.isen.com/index.html.

[9]. Felten, *supra* note 3, at 1.

[10]. *Id*.

This kind of design is not only more efficient, argues Felten, but would also encourage innovation at the edges of the network. In particular, Professors Lessig and Tim Wu of Columbia University Law School maintain that content providers will refrain from engaging in innovation if access providers have the ability to discriminate against some of them.[11]

This may sound reasonable on its face. But innovation among content providers does not appear to be slowing in spite of the lack of assurances that Lessig and company wish to legislate. The popularity of online search systems created by Yahoo! and Google has given way to upstart social networks like MySpace and YouTube and movie-delivery sites like Vongo.[12]

Much of the early writings on Transfer Communication Protocol/Internet Protocol (TCP/IP)—the technical rules of the Internet—is contained in a series of informal papers known as Requests for Comments.[13] These papers were not prepared by consultants on behalf of commercial interests. Instead, they functioned much like fodder for a chat room, offering design concepts for the Internet and applications of computer networking.

Interestingly, many of the contributors recognized the need to offer priority to some packets over others. A 1974 Request for Comments, whose authors included Google's current chief technologist who is now a major proponent of net neutrality, explained that incoming packets should be given priority over other packets on the edges of the network, to prevent congestion on the ingoing and outgoing pipe.[14] A 1981 Request for Comments explained that precedence—a measure of importance of the data stream—

---

[11]. Letter by Timothy Wu & Lawrence Lessig, Ex Parte Submissions, CS Dkt. No. 02-52, Aug. 22, 2003, at 4 ("The question an innovator, or venture capitalist, asks when deciding whether to develop some new Internet application is not just whether discrimination is occurring today, but whether restrictions might be imposed when the innovation is deployed.").

[12]. Nielsen/Netratings, Social Networking Sites Grow 47 Percent, Year Over Year, Reaching 45 Percent Of Web Users, May 11, 2006, *available at* http://www.nielsen-netratings.com/pr/pr_060511.pdf.

[13]. The first paper was published at UCLA in 1969. *See* 30 Years of RFCs, Apr. 7, 1999, *available at* ftp://ftp.rfc-editor.org/in-notes/rfc2555.txt.

[14]. Vinton Cerf, Yogen Dalai & Carl Sunshine, RFC 675-Specifications of Internet Transmission Control Program, Dec. 1974, *available at* http://www.faqs.org/rfcs/rfc675.html. This is not the only time that Vinton Cerf has made arguments in favor of prioritization. *See*, *e.g.*, Vinton Cerf, RFC 794-Pre-Emption. Sept. 1981, *available at* ftp://ftp.rfc-editor.org/in-notes/rfc794.txt. ("In packet switching systems, there is little or no storage in the transport system so that precedence has little impact on delay for processing a packet. However, when a packet switching system reaches saturation, it rejects offered traffic. Precedence can be used in saturated packet switched systems to sort traffic queued for entry into the system. In general, precedence is a tool for deciding how to allocate resources when systems are saturated. In circuit switched systems, the resource is circuits; in message switched systems the resource is the message switch processor; and in packet switching the resource is the packet switching system itself." It should be noted that Cerf's discussion of precedence here was in the context of some other, pre-Internet networking system. Still, it illustrates that prioritization is nothing new and has long been considered as a way to enhance a networking system.

could be used as a means of differentiating high-priority traffic from low-priority traffic.[15] A 1994 Request for Comments predicted that bandwidth constraints would eventually harm the delivery of real-time applications (think of live voice communication), and suggested that an arrangement for some traffic to receive preferred treatment was advisable.[16] These early writings on the Internet indicate that prioritization is not a new concept in Internet architecture. The end-to-end principle never was the sacrosanct design tenet that net neutrality proponents make it out to be.

Even if the end-to-end principle had been applied faithfully in the early stages of the Internet (which it was not) it is virtually irrelevant today. There are currently several ways suppliers of information on the Internet manage to get selected content and applications to users faster and more reliably. For their part, end-users can also buy services to improve speed and reliability. Modern networks of access providers like AT&T, Verizon, and Comcast can support quality-of-service (QoS) technology, which can give priority to specified traffic. During times of congestion, the lower priority traffic would be dropped first. QoS has multiple dimensions, including reliability, throughput, and speed.[17]

---

[15]. Information Sciences Institute, Internet Protocol Darpa Internet Program Protocol Specification, RFC 791, Sept. 1981, *available at* http://www.ietf.org/rfc/rfc0791.txt.

[16]. Robert Braden, David Clark & Scott Shenker, Integrated Services in the Internet Architecture: An Overview RFC 1633, Jun. 1994, at 3, *available at* http://www.ietf.org/rfc/rfc1633.txt.

[17]. See Robert E. Litan & Hal J. Singer, *Slouching Towards Mediocrity: Unintended Consequences of Net Neutrality Regulation*, Nov. 2006, *available at* http://ssrn.com/abstract=942043.

TABLE 1: QUALITY OF SERVICE OFFERINGS BY ACCESS PROVIDERS

ACCESS PROVIDER	SELECTED SERVICES & DESCRIPTIONS
AT&T	**Network-Based VPN** is a type of Virtual Private Network (VPN), a service that can be purchased by businesses with multiple locations to ensure high service levels for their traffic. It features advanced routing technology that prioritizes route assignments and quality of service and class of service traffic engineering capabilities for critical applications. **Premises-Based VPN** features "intelligent devices such as firewalls or VPN tunneling" and the ability to set appropriate levels of network security, user access control, and bandwidth prioritization. **Intelligent Content Distribution Service** is a service that can be purchased by content providers to ensure faster delivery of their data. According to the company, "AT&T leverages its own backbone, making AT&T uniquely positioned to find the best way to manage the traffic on content distribution… end-users can get the fastest possible download since content is placed on multiple servers located closer to your web viewers."
Broadwing	**Converged Services network** offers application-based class of service and quality of service options. "Customers can choose ports ranging in speed and may allocate bandwidth in the increments they require." **Remote Data Protection** is a "secure, network-based data protection service… fully automated remote data backup and recovery service for enterprises and their branch offices."
Global Crossing	**IP VPN Service**, through class of service and quality of service and associated service level agreements, allows customers to enjoy a high level of performance--"ideal for customers such as government, financial, media, and entertainment organizations."
Qwest	**Private Routed Networks (VPN)** is a service that offers higher security options for customers, such as protection from denial of service attacks that can cause web pages to be unavailable to the internet.
SAVVIS	**Managed IP VPN** is a service with "the ability to assign individual service levels to different applications so each application receives the performance levels required, and the customer pays for only what is needed." **Content Delivery Network**, a service that improves the performance and reliability of web applications, includes the following: caching services that enable the swift delivery of media content; streaming services that enable delivery of single events or libraries of video, music or animated content; intelligent traffic management software that routes "traffic to individual servers based on business rules" and reroutes it whenever performance bottlenecks emerge.

VERIZON	**IP VPN Dedicated** includes the following features for customers: traffic shaping or bandwidth allocation that provides real-time prioritization of outbound data and access control lists and a router-based firewall that provides a layer of traffic security, protecting against unwanted access to the customers' network.

Source: Company websites and 10-K filings with the Securities and Exchange Commission, accessed in September 2006; Litan and Singer (2006).

Table 1 shows that there are many different service offerings by access providers aimed at improving QoS. Other firms that are not access providers, including Akamai Technologies, CacheFly, Limelight Networks and Panther Express, offer similar services to content providers. Akamai, for example, provides content-acceleration service by caching content closer to the end-user for over 2,000 customers. There are also a large number of content providers, ranging from search engines like Google to businesses like eBay, that arrange to get their "products" to the market more quickly and reliably. Google sets up server farms packed with computers to store all of its content for end-users. Other firms are providing QoS for applications, like business-network alarm monitoring, that require a high degree of security.

By the same token, many applications depend on QoS to perform well. Popular QoS-needy applications include streaming multimedia, online gaming, Voice over Internet Protocol (VoIP), video teleconferencing, alarm signaling, and safety-critical applications such as remote surgery.

Online gaming provides a good example of how and why all bits are not treated equally by access providers. If there is even a small delay in response time with some games or degradation in the quality of the video stream, product quality declines unacceptably. The suppliers of these games will frequently pay web-hosting companies to offer faster and more reliable service than they could achieve with their own servers. They may pass the costs through to their customers. For example, users pay between $13 and $15 a month to subscribe to the popular multiplayer online role-playing game, World of Warcraft, part of which presumably goes to maintain the quality of the gaming network.[18]

Bit Equality and Net Neutrality Proposals

The net neutrality bills before the 109th Congress represented an attempt to regulate the *pricing* of service quality by an access provider. But, as is demonstrated in virtually all other sectors of the economy, pricing flexibility is generally a good thing.

Net neutrality proponents nonetheless argue that Internet access providers—cable companies, phone companies, media conglomerates—wield too much power over content providers, and as a result, cannot be trusted to charge fairly for service quality. Without such regulation, they say, access providers might one day monopolize the content markets by charging excessive prioritization fees to outside content companies. Moreover, if upstart content providers are protected from the possibility of price gouging, they will be encouraged to innovate faster than they do today. Net neutrality proponents assert that regulation before the fact (as opposed to antitrust suits after the fact) is necessary because the harm from

[18]. Rates are *available at* http://www.worldofwarcraft.com/info/faq/general.html.

anticompetitive behavior (in terms of less innovation by content providers) could not be remedied appropriately by either injunctive relief or fines. This line of argument discounts the counterargument—namely, that there is a possibility that such regulation could generate harm that far outweighs the potential benefits.

The net neutrality proposals would have achieved their objective by imposing non-discrimination requirements on access providers in the provision of QoS to content providers. Non-discrimination typically implies similar treatment for similar types of customers. For example, a non-discrimination rule for, say, a newspaper would require that a 2-inch-by-2-inch advertisement cost the same for all advertisers, regardless of the nature of the ad or its placement. Non-discrimination has a superficial appeal, but it is not always consistent with economic efficiency. Suppose that advertiser A was willing to pay 10 times more than advertiser B for an advertisement, but the newspaper is constrained to charge both advertisers the same price. Economic theory suggests that the total benefits to buyers and sellers can be increased by raising advertiser A's price while slightly lowering advertiser B's price—much the way air travelers who book at the last minute pay more than vacationers who book months ahead.[19]

Each of the net neutrality bills introduced in the 109th Congress, however, defined non-discrimination in the pricing of service quality as something more extreme: if a broadband provider offers enhanced service to any individual content provider, it must offer the same enhancements to all content providers at no extra charge. The idea is to stymie efforts by any content provider to secure enhanced service quality from access providers and instead to force all contracting for quality to occur between broadband providers and end-users.[20] Note, by the way, that these bills generally did not distinguish between broadband services offered by access providers and those offered by the companies that maintain the Internet "backbone," and they would have presumably imposed identical restrictions on both types of networks. Thus, because much of the enhanced QoS in the marketplace today resides at the backbone layer, those offerings would presumably be in jeopardy.

One bill, for example, known as H.R. 5273, sponsored by Representative Ed Markey, a leading Democrat on telecommunications, explained in its preamble that "a network-neutrality policy based upon the principle of nondiscrimination is essential to ensure that broadband telecommunications networks, including the Internet, remain open to independent service and content providers."[21] The bill would have required that an access provider "not discriminate in favor of itself in the allocation, use, or quality of

---

[19]. *See* JEAN-JACQUES LAFFONT & JEAN TIROLE, A THEORY OF INCENTIVES IN PROCUREMENT AND REGULATION 172-73 (MIT Press 1994).

[20]. *See*, *e.g.*, *Hearing on Net Neutrality Before the S. Comm. on Commerce, Science, and Transportation*, 109th Cong. (2006) (statement of Lawrence Lessig). AEI-Brookings Joint Center Testimony 06-01. March 2006, *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1254. ("To oppose access tiering [with content providers], however, is not to oppose all tiering. I believe, for example, that consumer-tiering should be encouraged. Network providers need incentives to build better broadband services. Consumer-tiering would provide those incentives.")

[21]. H.R. 5273, 109th Cong. § 2 (2006) [hereinafter H.R. 5273].

broadband services or interconnection with other broadband networks."[22] In other words, an access provider must offer the same service quality for its own content and unaffiliated content.[23] Finally, if an access provider offers a given service quality to one content provider, then it must offer the same service quality to all content providers free of charge.[24] Another net neutrality bill, S. 2360, sponsored by Senator Ron Wyden, an Oregon Democrat, would have similarly prevented an access provider from discriminating in the provision of QoS to content providers,[25] and it would ban any charges for QoS.[26] The bill would have also prohibited an access provider from discriminating with respect to bandwidth.[27] Although these bills were not passed in the 109th Congress, their prospects could be greatly improved in the current 110th Congress because of the Democratic majority. Already, one bill, known as S. 215, sponsored by Democratic Senator Byron Dorgan, has been introduced with similar language.[28]

Requiring that service quality be priced at zero for content providers could have dramatic effects on existing contracts between access and content providers, as those content providers would presumably seek to renegotiate their terms for service quality once the regulation had passed. Why should online gamers have to pay for QoS under existing contracts while everyone else gets the same service quality for free?

Net Neutrality is Likely to Produce More Harm Than Good

As the evidence on price implies, the market for broadband access is highly competitive. For example, between 2001 and 2006, the average price of a digital subscriber line dropped by about one-

---

22. *Id*. § 4(a)(5).

23. Access providers must ensure that unaffiliated content is delivered "at least equal to the speed and quality of service that the *operator's* content, applications, or service is accessed and offered, and without interference or surcharges on the basis of such content, applications, or services." *Id*. § 4(a)(6) (emphasis added).

24. "[I]f a broadband network provider prioritizes or offers enhanced quality of service to data of a particular type, [then it must] prioritize or offer enhanced quality of service to all data of that type (regardless of the origin of such data) *without imposing a surcharge* or other consideration for such prioritization or quality of service." *Id*. § 4(a)(7) (emphasis added).

25. S. 2360, 109th Cong. § 4(a)(6) (2006) (An access provider must "treat all data traveling over or on communications in a non-discriminatory way.").

26. *Id*. § 4(a)(4) (An access provider must "offer communications such that a subscriber can access, and a content provider can offer, unaffiliated content or applications or services in the same manner that content of the network operator is accessed and offered, without interference or surcharges.").

27. *Id*. § 4(a)(2)(A) (An access provider must "not discriminate in favor of itself or any other person, including any affiliate or company with which such operator has a business relationship in—(A) allocating bandwidth."). Another net neutrality bill, S. 2917, would have prevented an access provider from discriminating against a content provider with respect to bandwidth or QoS. *See* S. 2917, 109th Cong. § 12(a)(4)(A) (2006) [hereinafter S. 2917]. Access providers could offer prioritization to end-users but could not impose a fee for such service. *See id*. § 12(a)(5).

28. S. 215, 110th Cong. § 2 (2007) [hereinafter S. 215].

third.[29] In the case of cable, the quality-adjusted price declined significantly, as cable connection speeds increased significantly while prices held steady. Where there remains insufficient competition, the government's existing antitrust authority is sufficient to police an access provider's behavior. If, for example, a provider with monopoly power offered high quality service to an online gaming provider but refused to sell the same level of service to an unaffiliated voice-over-Internet protocol provider in order to protect its own subsidiary in the phone business, the antitrust laws would open the access provider to a lawsuit.[30]

But net neutrality would do more than gild the lily. Broadband access providers would likely react to non-discrimination provisions in ways that would hardly serve the interests of content providers and content users. An access provider could comply with a non-discrimination provision by either withdrawing its enhanced QoS offering from the marketplace or by replacing its tiered QoS offerings with a one-size-fits-all QoS plan. Under either scenario, the total benefits to consumers would diminish. To borrow another example from the airline industry, imagine a rule that required carriers to charge all customers the same price. One solution to this uniform pricing constraint would be for the airline to replace its first class and economy seats with some blended offering inferior to first class but superior to economy. To preserve revenues, the airline would be forced to set the price of the blended offering above the price of the economy seat. Customers who would have preferred to pay a lower price and receive a smaller seat would clearly suffer under such a regime, as would people who could no longer fly due to the higher prices. Moreover, business and first-class travelers who wanted better service and were willing to pay for it would have to settle for aching knees, pretzels instead of mixed nuts, and the occasional misbehaving child.

Entrepreneurs wishing to deliver content that depends critically on QoS would have no easy way to arrange it. Accordingly, they would likely divert their resources and creative energies to Internet uses that do not require high service quality. Consider the next generation of online video. Although today's video clips may not require exceptional bandwidth or reliability, as online video is narrowcast in high-definition format and takes on a more interactive nature, it is not much of a stretch to envision the day when content providers will be stymied by lack of resources. Thus, as a condition of investing resources into the development of online video, companies like Apple, Google, Vongo, and Sony need assurances that contracting for higher service quality with access providers will be legal.

---

[29]. Verizon charged $79 a month for a DSL connection in 2001. *See* Tom Spring, *Verizon Joins Broadband Price Hike Parade*, May 2, 2001, *available at* http://www.pcworld.com/resource/article/0,aid,48945,00.asp. In 2006, AT&T charged $30 a month. *See* Marguerite Reardon, *AT&T brings new low to DSL prices*, Feb. 3, 2006, *available at* http://news.com.com/AT38T+brings+new+low+to+DSL+prices/2100-1034_3-6034788.html.

[30]. A case in Canada mildly resembles this hypothetical fact pattern. In March 2006, Vonage Canada petitioned the Canadian Radio-television and Telecommunications Commission to investigate the conduct of Shaw Communications, a Canadian broadband cable modem service provider, of offering a monthly "VoIP tax" of $10. *See* Press Release, Vonage Canada, Who Controls How You Use Your Internet Access? Vonage Canada Challenges Shaw "VoIP tax" (Mar. 7, 2006).

Future welfare depends on innovation by both access providers *and* content providers. By eliminating or seriously jeopardizing the market for real-time applications, net neutrality legislation would likely reduce consumer welfare for current real-time applications, such as online gaming. It would also dampen the development of new applications that depend critically on quality of service.

Conclusion

Although it may have represented a democratic ideal two decades ago, the end-to-end principle is a fiction today and should be treated as such. Policymakers should look at how the Internet really functions from both a technical and an economic perspective. Modern networks are capable of acting more intelligently than earlier networks, which means that access providers now have the ability to prioritize data for those applications that depend critically on it. If content providers are willing to pay for enhanced quality, there is no good reason for regulators to deter them.

What will protect upstart content providers in the absence of net neutrality regulation? Two things: The first, as noted earlier, is antitrust laws that deter access providers with demonstrated market power from abusing their positions. The second is a competitive environment in which the prospect of gaining market power over broadband access is remote. Along with 25 other economists, from William Baumol to Alfred Kahn to Hal Varian, we recently signed a statement calling for an end to local franchising regulations for broadband and the federal transfer of more wireless spectrum to private broadband uses.[31] And, we would argue, it is here—and only here—that there is room for government regulation that is likely to do more good than harm.

---

[31]. *See* Elizabeth E. Bailey, Martin Neil Baily, William J. Baumol, Peter Cramton, Gerald R. Faulhaber, Kenneth Flamm, Richard Gilbert, Austan Goolsbee, Shane Greenstein, Robert W. Hahn, Robert E. Hall, Thomas W. Hazlett, Alfred E. Kahn, Robert E. Litan, John Mayo, Paul Milgrom, Janusz A. Ordover, Robert S. Pindyck, Gregory L. Rosston, Scott J. Savage, Howard Shelanski, Richard L. Schmalensee, Pablo T. Spiller, David J. Teece, Hal R. Varian, Scott Wallsten, Dennis L. Weisman, *Economists' Statement on U.S. Broadband Policy*, AEI-Brookings Joint Center Related Publication 06-06, Mar. 2006, *available at* http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1252.