



J O I N T C E N T E R
AEI-BROOKINGS JOINT CENTER FOR REGULATORY STUDIES

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

Robert W. Hahn and Robert E. Litan*

Related Publication 06-33
November 2006

This paper can be downloaded free of charge from the AEI-Brookings Joint Center's website
www.aei-brookings.org or from the Social Science Research Network at:
<http://ssrn.com/abstract=947847>

A revised version of this paper is forthcoming in the *Milken Institute Review*.

* Mr. Hahn and Mr. Litan co-founded and direct the AEI-Brookings Joint Center. They have consulted for telecommunications companies and information technology companies on issues discussed in this paper. The authors wish to thank Caroline Cecot, Hal Singer, Scott Wallsten and Molly Wells for helpful suggestions. The views expressed here are their own.

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.

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

Robert W. Hahn and Robert E. Litan

I. Introduction

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, written by MIT computer scientists Jerome Saltzer, David Clark and David Reed, extolled a design principle called “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, 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 equally, and letting the intelligence reside principally at the edges of the network, entrepreneurs can compete to bring consumers new products or applications, or so the argument goes.

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 for two reasons. First, access providers are in the midst of a multi-billion 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, access providers 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. Second, the

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>.

prospects for legislation have improved because key Democratic supporters are back in leadership positions.

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 show that because 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, though we recognize many net neutrality advocates—who frequently support treating all bits equally in terms of not discriminating—would disagree. Third, we consider our findings in the context of the current legislative proposals involving net neutrality. We argue that these proposals are misguided and extreme—even by the standards of those advocating net neutrality. In particular, the proposals would require an access provider that offered enhanced quality of service to one content provider to make the same level of service quality available to all content providers free of 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 generally a good thing, and there is no reason to believe that pricing flexibility is not a good thing here. We argue that the antitrust laws are sufficient to police access providers with market power that discriminate among unaffiliated content providers in the provision of service quality.⁴

II. 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 connected by long wires. Packets of data

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, *Telecommunications, the Transition from Regulation to Antitrust*, AEI-Brookings Joint Center Related Publication 06-21, Jul. 2006, available at <http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1314>; 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 Owen, Bruce 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>.

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. The speeds are 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 advertised maximum speeds. Thus, a consumer may pay \$15 per month for a maximum speed of 786 Kbps and \$30 for maximum speeds up to 3 Mbps.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 prioritization and 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 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 those features a second time on an end-to-end basis.⁷ This viewpoint leads to the model of “dumb network”⁸ with intelligence built into the edges of the network only. According to Edward Felten, a proponent of the end-to-end principle, the routers in the middle of the Internet currently “forward packets with only minor processing—all the heavy lifting takes place on the transmitting and receiving computers.”⁹ Felten argues for keeping things this way in large part because computers at the edges of the network are owned and controlled directly by end users.¹⁰

Proponents of the end-to-end principle suggest that it is not only more efficient for the reasons identified by Felten, but would also encourage innovation at the edges of the network. For example, in their joint letter to Congress in August 2003, Lawrence Lessig and Timothy Wu

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://www2.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.

claimed that, by failing to ensure the principle of end-to-end through regulation, the FCC and Congress harm future innovation by causing uncertainty among innovators and entrepreneurs.¹¹ In particular, they argue that content providers will refrain from engaging in innovative activities due to the fact that access providers have the ability to discriminate against content providers on the basis of type of application.

This may sound reasonable, at first blush, but it misses the fact that innovation among content providers does not appear to be slowing. The popularity of online search algorithms created by Yahoo and Google has given way to upstart social networks like MySpace and YouTube.¹² Those who attribute the boom in innovation solely to the Internet's end-to-end nature have not made their case. The Internet is not end-to-end now and was never designed to be strictly neutral.

Most of the early writings on Transfer Communication Protocol/Internet Protocol (TCP/IP) – the technical rules of the Internet -- are 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, which happened to be co-authored by Google's current chief technologist and, ironically, a major proponent of net neutrality today, explained that outgoing packets should be given priority over other packets to prevent congestion on the ingoing and outgoing pipe.¹⁴ A 1981 Request for Comments explained that precedence—a

¹⁰ *Id.*

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

measure of importance of the data stream—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 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 has always been considered an important design characteristic for TCP/IP -- sharp contrast to the romantic ideal of the end-to-end principle.

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. In addition, there are several ways users can receive those services faster and more reliably. Modern networks of access providers support quality-of-service (QoS) technology, which can label some traffic as higher priority than other traffic. During times of congestion, the lower priority traffic would be dropped first. QoS has multiple dimensions, including reliability, throughput, and speed.¹⁷

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.”

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	<p>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.</p> <p>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.</p> <p>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.”</p>
Broadwing	<p>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.”</p> <p>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.”</p>
Global Crossing	<p>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.”</p>
Qwest	<p>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.</p>
SAVVIS	<p>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.”</p> <p>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.</p>
VERIZON	<p>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.</p>

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, CacheFly, Limelight Networks and Panther Express, offer similar services for content providers. Akamai Technologies 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, who make arrangements 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 such as alarm monitoring that require a high degree of security.

There are many applications that depend on QoS to perform properly. 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. The quality of a VoIP stream is highly sensitive to both time delay and packets that arrive out of order.

Online gaming provides a good example of how and why all bits are not treated equal 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 per 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

III. Bit Equality and Net Neutrality Proposals

The net neutrality bills before Congress represent 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 access providers wield too much power over content providers when it comes to delivering content to end users, and as a result, these providers cannot be trusted to negotiate fairly for the pricing of service quality. Without such regulation, an access provider might one day attempt to monopolize the content markets by charging excessive prioritization fees to upstart content providers. If upstart content providers are protected by regulators from those surcharges, the argument goes, then they will be encouraged to innovate faster than they do today. Net neutrality proponents argue that *ex ante* regulation is necessary because the harm from anticompetitive behavior (in terms of less innovation by content providers) could not be remedied appropriately by antitrust courts by either injunctive

¹⁸ Rates are available at <http://www.worldofwarcraft.com/info/faq/general.html>.

relief or monetary fines. This line of argument discounts or ignores the fact that *ex ante* regulation can generate other harms, as described in the following section, that exceed the alleged harms that ostensibly justify net neutrality regulation.

The net neutrality proposals would achieve 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. Non-discrimination has a superficial appeal, but it is not always consistent with economic efficiency. Suppose that advertiser *A* was willing to pay ten times more than advertiser *B* for a 2-by-2 inch 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 welfare is increased by charging air travelers who book at the last minute more than vacationers who book months ahead.¹⁹

Under each of the net neutrality bills in Congress, however, non-discrimination in the pricing of service quality means something more extreme: If a broadband provider offers enhanced service quality to any individual content provider, then it must offer the same enhanced level of service quality to all content providers for free. The apparent motivation for such a restriction is to stymie efforts by any content provider to secure enhanced service quality from access providers, and instead to force all contracting for service quality to occur between broadband providers and end users.²⁰ These bills generally do not distinguish between broadband services offered by access providers versus those offered by backbone networks, and they would presumably impose their net neutrality restrictions on both types of networks. Because much of the enhanced QoS in the marketplace today, including local caching of content, resides at the backbone layers of network, those offerings would presumably be in jeopardy.

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.”)

A review of the key language from the net neutrality proposals in Congress reveals just how far the net neutrality proponents are willing to go to achieve bit equality. To understand why the net neutrality proposals represent bad economics, one must first review the key provisions in these bills relating to service quality. One net neutrality bill in the House, H.R. 5273, explains 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 With respect to content providers, the bill would require that access providers “not discriminate in favor of itself in the allocation, use, or quality of broadband services or interconnection with other broadband networks.”²² 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, similarly would prevent an access provider from discriminating in the provision of QoS to content providers,²⁵ and it would ban any charges for QoS.²⁶ S. 2360 also would deny an access provider from discriminating against either a content provider or end user with respect to bandwidth.²⁷ Another net neutrality bill, S. 2917, would prevent an access provider from discriminating against a content provider with respect to bandwidth or QoS.²⁸ Access providers could offer prioritization to end users but could not impose a fee for such service.²⁹

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

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.”).

²⁸ S. 2917, 109th Cong. § 12(a)(4)(A) (2006) [hereinafter S. 2917].

²⁹ *Id.* § 12(a)(5).

In summary, under each of the net neutrality proposals, access providers would be severely limited in their offerings to content providers and moderately limited in their offerings to end users. Requiring that service quality be priced at zero for content providers could have dramatic effects on existing contracts between access providers 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?

IV. Net Neutrality Is Likely To Produce More Harm Than Good

For at least three reasons, net neutrality regulation is not needed and would likely reduce consumer welfare. First, it is not clear how proactive regulation would generate social benefits in excess of social costs in the presence of antitrust law. Broadband access markets are highly competitive as evidenced by rapidly declining prices and entry by wireless providers. Since 1999, the price of a DSL connection has declined from roughly \$80 per month to less than \$20 per month. While the absolute price of a cable modem has not declined as rapidly, the quality-adjusted price has declined significantly, as cable modem connection speeds have more than doubled while prices held steady. Where there remains insufficient competition, the government's existing antitrust authority is a sufficient tool to police an access provider's behavior. Suppose that a hypothetical monopoly access provider offers high QoS to an online gaming provider but refuses to sell the same level of QoS to an unaffiliated VoIP provider to protect its affiliated VoIP provider.³⁰ Antitrust laws police such behavior by not permitting such discriminatory refusals to deal among unaffiliated content providers.

Second, access providers would likely react to the non-discrimination provisions in ways that are harmful for content providers and their end users. An access provider may attempt to comply with a non-discrimination provision in the supply of QoS by either withdrawing its enhanced QoS offering from the marketplace or by replacing its tiered QoS offerings with a one-size-fits-all QoS offering. Under either scenario, consumer welfare associated with the purchase of enhanced QoS would be greatly reduced. To borrow an example from the airline industry,

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

consider a rule that required airlines 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, which was inferior to first class but superior to traditional economy. To preserve average revenue per customer, 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 can no longer fly due to the higher prices. Moreover, business and first-class travelers who place a high value on sitting in a preferred class would also experience lower utility under such a regime by having to sit with chatty tourists (and the occasional misbehaving child).

Third, content providers interested in designing and producing content that depends critically on QoS would have no means of providing that content, at least not in an acceptable manner. Accordingly, they would likely divert their resources and creative energies to other applications that do not require high service quality. Consider the next generation of online video. Although current video clips may not require high service quality (guaranteed throughput may be required for streaming video), as online video takes on a more interactive nature, it is not much of a stretch to envision how Apple or some other video provider would demand high service quality from access providers. As a condition of investing resources into the development of online video, companies such as Apple and Sony need assurances that contracting for higher service quality with access providers will be legal.

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.

V. Conclusion

Although it may have represented a democratic, romantic 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 the antitrust laws that deter access providers with demonstrated market power from abusing their position. The second is a competitive environment in which the prospect of gaining market power over broadband access is remote. Along with 25 other economists, 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 SSRN: <http://ssrn.com/abstract=892009>.