



ISSUE BRIEF



FEBRUARY 2014

PUBLICPOLICY.WHARTON.UPENN.EDU

VOLUME 2, NUMBER 2

UPDATING INTERNET POLICY FOR THE 21ST CENTURY

CHRISTOPHER S. YOO

January 14 of this year will go down as a landmark in the history of the Internet. On that day, the U.S. Court of Appeals for the D.C. Circuit invalidated the Open Internet Order, which was the Federal Communications Commission's (FCC's) effort to mandate network neutrality.

The Order would have prevented broadband access providers, such as Comcast and Verizon, from charging applications or content providers for prioritized service. In short, network neutrality stands for the principle that all bits are created equal and should be treated the same.

The court's opinion is complex and highly legalistic, giving both sides some things to celebrate and some things to lament. Rather than parse the finer points of the court's reasoning, I would like to situate the decision in its broader historical and technological context. The idea of network neutrality has its roots in the simpler times of the mid-1990s, when a small number of academics and technophiles used a personal computer (PC) connected to a telephone line to send email and browse the web.

But it is no longer suitable as a governing principle for Internet policy in the 21st century. As I point out in my most recent book, *The Dynamic Internet*, the network has changed a great deal over time. Today, the Internet has become much larger and more diverse in terms of *users, applications, technologies*, and *business relationships*. These changes have raised doubts about the one-size-fits-all approach reflected in network neutrality and created pressure to allow different actors to experiment with a broader variety of solutions that deviate from those of the past. In this Issue Brief, I will outline the major changes the Internet has undergone in recent years and discuss their implications for the future of Internet regulation.

ABOUT THE AUTHOR



Christopher S. Yoo

Professor of Law, Communication, and Computer & Information Science, Penn Law School

Christopher S. Yoo is the John H. Chestnut Professor of Law, Communication, and Computer & Information Science and Founding Director of the Center for Technology, Innovation & Competition at the University of Pennsylvania Law School. He has emerged as one of the nation's leading authorities on law and technology. His research focuses on how the principles of network engineering and the economics of imperfect competition can provide insights into the regulation of electronic communications. He has been a leading voice in the "network neutrality" debate that has dominated Internet policy over the past several years. He is also pursuing research on copyright theory as well as the history of presidential power. He is the author of *The Dynamic Internet: How Technology, Users, and Businesses Are Transforming the Network* (AEI Press, 2012), *Networks in Telecommunications: Economics and Law* (Cambridge Univ. Press, 2009) (with Daniel F. Spulber) and *The Unitary Executive: Presidential Power from Washington to Bush* (Yale Univ. Press, 2008) (with Steven G. Calabresi). Professor Yoo testifies frequently before Congress, the Federal Communications Commission, and the Federal Trade Commission.

THE INCREASE IN THE NUMBER OF END USERS

One of the biggest changes over the past two decades has been the explosion in the number of people using the Internet. The dramatic increase in the number of end users also necessarily means that they are more geographically dispersed as well as more diverse in terms of backgrounds and what they value about the Internet.

The tremendous rise in the number of Internet users has changed the way users interact with each other. During the early days of the Internet, the community relied on shame and peer pressure to prevent people from engaging in undesirable behavior, such as sending spam. Those days are clearly long gone. Now, the increase in the number of end users has made it impossible to rely on common values and informal sanctions to keep order. Moreover, the increase in the size in the Internet has been accompanied by a marked decrease in the level of trust between users. And in light of technological changes described below, it also means that more people are using the Internet in new and more varied ways that were not part of the picture when the concept of network neutrality was first articulated.

THE EMERGENCE OF INTERNET VIDEO

The applications that characterized the early days of the Internet were relatively simple. Email and web browsing, the applications that dominated the early Internet, did not use significant amounts of bandwidth. In addition, delays of up to half a second were often unnoticeable and certainly did not render the service unusable. Moreover, email and web browsing were not particularly sensitive to irregularities in the flow of packets (known as jitter), which could arise as a result of packet loss or congestion.

The modern Internet is dominated by applications that are much more demanding. One of the most prominent of these is video, provided by companies such as Netflix, Hulu, and Amazon. As an initial matter, video requires substantially more bandwidth than web browsing or email. Indeed, industry ana-


lysts report that Netflix by itself accounts for more than one-third of all primetime Internet traffic. Video is also very sensitive to delay and jitter. Half-second delays can cause the screen to lock up. If this happens too frequently, consumers will simply stop using the service.

Interestingly, for prerecorded video, these problems can be largely eliminated simply by delaying playback for a few seconds until a sufficient number of packets

100 Mbps service to 100 million homes could cost up to \$400 billion. Needless to say, in the aftermath of the economic downturn of 2008 and in a climate where the government is looking for ways to reduce spending, options that that would reduce the need to undertake such large capital expenditures need to be considered seriously.

Moreover, on a more technical level, the basic approach to managing congestion on

FIGURE 1: GLOBAL INTERNET USERS (BILLIONS), 1995-2013



Source: Internet World Stats (n.d.), "Internet Growth Statistics," <http://www.internetworldstats.com/emarketing.htm>.

are placed in temporary storage. Building up a buffer of packets allows the application to cushion the playback against any irregularities and to release the packets in a steady stream even if their pattern of arrival was more erratic. Buffering does not work for interactive video, such as video conferencing and some online gaming, which cannot tolerate the latency that arises when an application buffers packets.

The result is that many providers are attempting to deal with increases in the amount of traffic either by prioritizing video or by reserving bandwidth specifically for video, while giving lower priority to traffic that is less sensitive to delay. Although many would argue that this would deviate from the approaches used in the past under the principle of network neutrality, it seems to be a necessary change. The only alternative would be to add more capacity, but some estimate that providing

the Internet does not work for the transport protocol employed by most video applications to transport data (known as User Datagram Protocol or UDP). Unlike the transport protocol that dominated the early Internet, UDP does not back off when confronted with congestion. Although there have been some attempts to integrate UDP-based applications into existing approaches to congestion management, to date these efforts have not been wholly successful. This means that broadband networks may have to treat video packets differently from other packets. Broadband access networks may also have to engage in more extensive network management as the amount of video increases. The old ways of doing things are simply not practical anymore.

The growing importance of Internet video has also fundamentally altered the competitive dynamics of the industry and, in the process, has called into question some

of the assumptions underlying network neutrality. Network neutrality is animated by the concern that broadband access providers would use their market position to place economic pressure on content and application providers. Some network neutrality advocates have argued that regulation should foreclose network providers from exercising their bargaining power by mandating that the price that ISPs can charge content and application providers always be zero. More recently, however, the shoe has sometimes been on the other foot. Leading content providers, such as Netflix and ESPN, have been using the leverage created by their popularity to seek better commercial deals from broadband access providers. I do not mean to suggest that there is anything wrong with their doing so. To the extent that bargaining power is the result of financial risks and investments undertaken by each firm, this give-and-take is a normal part of a healthy economic market. My point is that it is a mistake to build policy around preconceived notions of the distribution of economic power, since the competitive dynamics are in constant flux and any presumptions about which side has the stronger bargaining position are very likely to shift over time. Indeed, although Netflix enjoyed considerable initial success in requiring network providers to terminate its traffic for free, the recent deal between Netflix and Comcast suggests that the pendulum may be swinging the other way.

THE RISE OF WIRELESS BROADBAND

Another major development is the growing importance of wireless technologies. In a few short years, wireless broadband has gone from having no subscribers to surpassing both cable modem and DSL as the leading platform for broadband services. Figure 2 shows the FCC's data for its benchmark service of 3 Mbps downstream and 768 kbps upstream, where mobile broadband represents 50% of all subscriptions, compared to 34% for cable modem service and 10% for ADSL.


If one looks at basic broadband of 200 kbps, wireless broadband becomes even more dominant, representing 65% of all sub-

scriptions as compared with 20% for cable modem service and 12% for ADSL.

If anything, Figures 2 and 3 understate the current importance of wireless broadband. The key development is the high-speed fourth-generation (4G) wireless technology known as long-term evolution (LTE). As of the end of 2012, Verizon's LTE network reached 87% of the U.S. population, with AT&T reaching 48%,


available, AT&T's, Verizon's, and T-Mobile's LTE networks are currently delivering average download speeds of 12 Mbps and peak download speeds of 60 Mbps or more. These speeds meet or exceed the recommended bandwidth requirements for Netflix (8 Mbps) and for multi-person video conferencing on Skype (12 Mbps). The future holds even more promise. Wireless providers in the UK, Korea, and Australia are already deploying upgraded versions

FIGURE 2: BROADBAND SUBSCRIPTIONS (3 MBPS) (MILLIONS), DEC 31, 2008 - DEC 31, 2012



Source: FCC (2013), "Internet Access Services: Status as of December 31, 2012," http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-324884A1.pdf.

FIGURE 3: BROADBAND SUBSCRIPTIONS (200 KBPS) (MILLIONS), DEC 31, 2008 - DEC 31, 2012



Source: FCC (2013), "Internet Access Services: Status as of December 31, 2012," http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-324884A1.pdf.

Sprint reaching 38%, and T-Mobile not yet having begun to deploy. By the end of 2013, Verizon had completed its LTE buildout, and AT&T reached 85%, T-Mobile reached 71%, and Sprint reached 63%. All four companies are projected to complete their buildouts by the middle of 2014.

The deployment of LTE should substantially increase wireless bandwidth. Where it is

of LTE capable of delivering download speeds of 150 Mbps and even 300 Mbps.


The business environment surrounding wireless broadband is starkly different from the business environment associated with the wireline Internet. As of the end of 2012, the FCC reports that 97% of the U.S. population lived in census tracts served by three or more providers offering service at the FCC's

benchmark level of 3 Mbps stream and 768 kbps upstream. The FCC cautions that these statistics may overstate the level of competition. This is because an entire census tract is considered covered by a provider so long as that provider serves a single household within that tract, even if that provider does not serve the entire tract. Nonetheless, the trend towards increasing competition is unmistakable.

The markets remain quite competitive at higher speed tiers. At the 6 Mbps downstream/1.5 Mbps upstream bench-

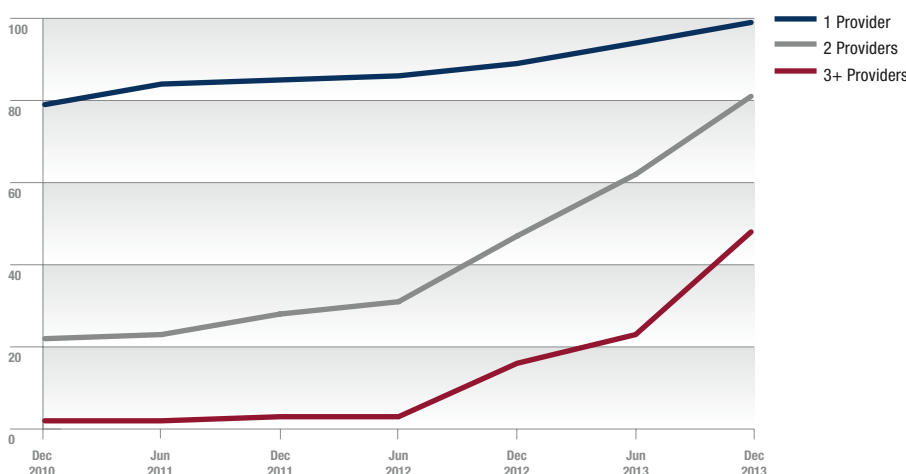
mark, 81% of the U.S. population lived in census tracts served by three or more providers. Even at the 10 Mbps downstream/1.5 Mbps upstream benchmark, which is the highest speed tier for which the FCC collects data, 48% of the U.S. population lived in census tracts served by three or more providers, and 80% lived in census tracts served by two or more providers. The broader deployment of LTE since that time has no doubt caused these numbers to rise still further.

FIGURE 4: PERCENTAGE OF U.S. HOUSEHOLDS LOCATED IN CENSUS TRACTS WHERE BROADBAND PROVIDERS OFFER DOWNLOAD SPEEDS OF 3 MBPS, DEC 31, 2008 - DEC 31, 2012



Source: FCC (2010-2013), "Internet Access Services," Status as of December 31, 2009, through December 31, 2012, <http://transition.fcc.gov/wcb/iatd/comp.html>.

FIGURE 5: PERCENTAGE OF U.S. HOUSEHOLDS LOCATED IN CENSUS TRACTS WHERE BROADBAND PROVIDERS OFFER DOWNLOAD SPEEDS OF 10 MBPS, DEC 31, 2008 - DEC 31, 2012



Source: FCC (2010-2013), "Internet Access Services," Status as of December 31, 2009, through December 31, 2012, <http://transition.fcc.gov/wcb/iatd/comp.html>.

As was the case with video, the emergence of wireless broadband has changed the focus of competition policy. Historically, the concern has been that broadband access providers would be able to exert market power against other parts of the industry. In the current environment, competition authorities have become just as concerned that manufacturers of leading wireless devices, such as the Apple iPhone and Google's Android-based phones, may be in a position to exercise market power in the other direction. Again, to the extent that the bargaining power enjoyed by any of these parties is the result of business acumen or foresight, preserving incentives to innovate requires that they be allowed to enjoy the fruits of their labors and willingness to take risk. The dynamic nature of the industry cautions strongly against basing policy on any presumptions about the sources of bargaining leverage.

End users also appear to use wireless broadband connections in ways that are fundamentally different from wireline connections. Instead of consuming different types of content located through a search engine, wireless users tend to focus on apps, which they find through the app store. This means that the relevant platform has shifted from browsers to wireless operating systems, such as Apple's iOS, Google's Android, or Microsoft's Windows Phone. At the same time, wireless users appear to be more willing to pay for apps than wireline users were willing to pay for content. Furthermore, the industry is experimenting with a wide range of new configurations, incorporating some functions normally considered applications into the operating system (e.g., Apple FaceTime) and others into the chip itself (e.g., Google Wallet). The net result is that the value chain in the wireless world is completely different from the value chain of the wireline world, with different sets of relative winners and losers.

Moreover, the technical environment associated with wireless broadband is far different from that of the wireline world dominated by cable modem service and ADSL. The primary source of these differences is the fact that wireless networks are much less reliable than wireline networks. This in turn requires wireless to deploy network-based error recovery techniques such as Automatic Repeat reQuests (ARQ). The problem is that

ARQ uses deep packet inspection (DPI) and a variety of other functions and embeds them deep within the network. The result is that broadband access providers necessarily must manage wireless broadband networks far more extensively and intrusively than was necessary for cable modem or DSL service. Again, this is a technical change that requires rethinking the old ways of regulating Internet activity.

THE MATURATION OF THE U.S. BROADBAND MARKET

When the broadband market was growing rapidly, providers had the incentive to offer a standardized product designed to draw in new customers. In recent years, however, the broadband market has approached saturation, with subscriber growth slowing dramatically.

This shift has caused the nature of competition to shift from *extensive* competition, in which firms seek to serve new customers who are entering the market, to *intensive* competition, in which firms seek to deliver higher value to customers who are already in the market. When competition shifts from extensive to intensive, the natural response is for providers to offer increasingly specialized services in an attempt to deliver more tailored services that individual consumers value more highly. In the context of the Internet, this may lead to greater use of the types of prioritized services that network neutrality is designed to prevent. It may also lead to firms making greater use of strategic partnerships and vertical integration. Policymakers must keep in mind that product differentiation can represent an important source of competitive rivalry and can provide real value to consumers, and that market-driven consolidation does not necessarily harm consumers.


THE MYTH OF THE ONE SCREEN

If end users maintain multiple connections, so long as they can access the content they desire, it should make no difference which connection they use to do so. The policy underlying network neutrality—that every connection should provide access to every website on equal terms—is based on the implicit presumption that every person will

subscribe to only one broadband service. Only if that is the case must every connection be everything to everyone.

A casual examination of people's actual behavior reveals a more complex outcome. As suggested by the data in Figures 2 and 3, most Americans subscribe to both a fixed line and a wireless broadband provider, largely because of their different technical characteristics.

FIGURE 6: PERCENTAGE OF U.S. ADULTS WHO ACCESS THE INTERNET VIA BROADBAND, MARCH 2000 - MAY 2013



Source: Pew Research Center (2013), "Home Broadband Adoption," [http://pewinternet.org/Static-Pages/Trend-Data-\(Adults\)/Home-Broadband-Adoption.aspx](http://pewinternet.org/Static-Pages/Trend-Data-(Adults)/Home-Broadband-Adoption.aspx).

Fixed-line services provide greater bandwidth. Wireless services provide mobility. In addition, many households continue to subscribe to cable, satellite, or some other form of multi-channel video. Still others rely on other firms for functions such as alarm monitoring.

The existence of multiple connections (a practice known as multihoming) weakens the leverage that any one broadband provider has over subscribers. This in turn allows policymakers to rely more on competitive dynamics, and less on regulation, to protect consumers. At the same time, it undercuts claims that every connection must meet the needs of every person. Instead, it opens up the possibility of different providers targeting their offerings towards different populations.

CLOUD COMPUTING

Cloud computing represents one of the most controversial developments regarding the

Internet, with many regarding it as a fundamental change in the architecture and others critiquing it as overhyped repackaging of existing technologies. During the PC era, individual end users relied primarily on the resources located in their desktop or laptop computers. In this environment, applications such as email or word processing relied on the CPU located in their PC for computing

power and relied on the hard drive located in the PC to store both the software and the data associated with the application.

Cloud computing applications, such as Gmail and Google Apps, follow a different model. Instead of relying on resources contained in the PC, cloud applications rely on computing power and storage facilities located in remote data centers maintained by the cloud provider. End users only need what are often called "thin clients," that is, very simple computers only sophisticated enough to run a browser that users can use to access cloud resources.

The rise of the cloud is in the process of rearranging the structure of the industry. Simplifying the software required on PCs and laptops has weakened the centrality of PC operating systems, such as Microsoft Windows. At the same time, it has heightened the importance of other economic actors. One example is VMware, which creates systems

known as hypervisors that organize and manage functions within data centers. Software producers must also stop thinking about software as a product, along with the attendant focus on periodic new versions made available on major release dates. Instead, they must think of software as a service characterized by an environment of constant improvement

Furthermore, the shift to the cloud requires that data that used to reside exclusively within an end user's PC must now pass through a network and reach a data center. This means that network connectivity must be ubiquitous for cloud solutions to work. Moreover, if the network is slow or unreliable, end users will find their cloud applications to be unusable. This may lead cloud users to insist on certain guaranteed levels of quality of service from their network providers. In addition, the fact that previously private information must pass through the network and share space on a server with other users may mean that cloud customers may begin to demand higher levels of privacy and security. This has led many initiatives to explore ways to redesign the Internet's architecture to permit prioritization of traffic and to improve

identify verification. Both changes could well require some deviations from the traditional vision of network neutrality.

CONCLUSION

In short, the Internet is now characterized by an economic and technological reality that is more complex than the one that existed in the mid-1990s. The natural response is for the industry to adapt to these changes by providing services that are more diverse. Although these innovations represent deviations from the status quo, they should not reflexively be regarded as harmful to consumers. Nor is there any reason to assume that the pace of change will slacken any time in the foreseeable future.

At the same time, because theoretically the changes that the Internet is undergoing could both benefit and harm consumers, some level of regulatory oversight is required. I have long advocated creating a regulatory regime based on case-by-case adjudication. This regime should intervene only if real-world data shows harm to consumers and places the burden of proof on the party challenging the practice. Any other approach

would make "no" the default response rather than "yes," thereby depriving innovation of the breathing room it needs to experiment with new solutions to new problems.

BRIEF IN BRIEF

- Over time, the Internet has become much larger and more diverse in terms of users, applications, technologies, and business relationships.
- These changes have called into question the idea of network neutrality (the principle that Internet service providers and governments should treat all data equally), which has shaped Internet policy since the 1990s.
- Moreover, with the emergence of wireless broadband and the growing importance of Internet video, it has become clear that the animus for network neutrality—the concern that broadband access providers would use their market power to control Internet content and application providers—does not fully reflect the true competitive dynamics of the industry.
- While some regulatory oversight of the Internet is required, the regulatory regime should be based not on the one-size-fits-all approach of network neutrality, but rather around case-by-case adjudication, whereby action is taken only when harm to consumers can be proven with real-world data.

Founded in 1881 as the first collegiate business school, the Wharton School of the University of Pennsylvania is recognized globally for intellectual leadership and ongoing innovation across every major discipline of business education. With a broad global community and one of the most published business school faculties, Wharton creates economic and social value around the world.

KNOWLEDGE FOR POLICY IMPACT

ABOUT THE PENN WHARTON PUBLIC POLICY INITIATIVE

The Penn Wharton Public Policy Initiative (PPI) is a hub for research and education, engaging faculty and students across University of Pennsylvania and reaching government decision-makers through independent, practical, timely, and nonpartisan policy briefs. With offices both at Penn and in Washington, DC, the initiative provides comprehensive research, coverage, and analysis, anticipating key policy issues on the horizon. Penn Wharton PPI is led by Faculty Director Mark Duggan, the Rowan Family Foundation Professor and Chair of Business Economics and Public Policy, and Professor of Health Care Management at Wharton.

ABOUT PENN WHARTON PPI ISSUE BRIEFS

Penn Wharton PPI publishes issue briefs at least once a month, tackling issues that are varied but share one common thread: they are central to the economic health of the nation and the American people. These Issue Briefs are nonpartisan, knowledge-driven documents written by Wharton and Penn faculty in their specific areas of expertise.

For additional copies, please visit the Penn Wharton PPI website at publicpolicy.wharton.upenn.edu.

 Follow us on Twitter: @PennWhartonPPI