Telecommunications:
Net Neutrality: Battle of the Titans

William P. Barr, Paul Misener, Christopher S. Yoo; Moderator: David M. McIntosh

Introduction: I'm David McIntosh, a partner at Mayer Brown. I'll be the moderator for today's panel on Net Neutrality and the future of that issue. I want to give you a little bit of background. As you know, in the 109th Congress, there was a large, long effort to bring telecom deregulation bills forward. Towards the end of that effort last year, the issue of Net neutrality surfaced as the major stumbling block for that deregulatory effort. In the House, there were amendments to bring Net neutrality as part of the Bill that were defeated. In the Senate, it was still being negotiated, and incrementally different versions of Net neutrality were included in the Bill that Senator Stevens had worked in his commerce committee.

Now although last I heard, Senator Stevens still hold out some hope that in this lame-duck session there might be a compromise on a bill passed, most people feel that it's very unlikely and that the issue will be returning once again in the 110th Congress, only this time with very different committee leadership. Chairman John Dingell, whom I served with, is going to be Chairman of the Commerce Committee. Ed Markey, who is a big proponent of Net neutrality, will chair the Telecommunications Subcommittee in the House. In the Senate, Senator Inoue will take over chairing the Commerce Committee. His staff has indicated that they want to take a look at Net neutrality in a serious way. There will also be changes, I think, that come ultimately as a result of last Tuesday's elections, on the way issues are decided in other areas of government. Chairman Dingell has indicated that he thinks the '96 Act is outdated and wants to bring forward some type of telecom act. But it won't be the same as it was in the 109th Congress. He has plans to modify the universal service fund and look at media ownership.

One thing that I want to mention at the outside is that the term "Net neutrality" has different meanings to different people. There's the minimalist approach, as I call it, which says essentially that there should be limits on Internet service providers, stopping them from prohibiting any of their subscribers from posting or delivering in e-mail or trap visiting any site on the web. Then, provisions from the Markey Bill that say, in addition to not denying access, providers should be prohibited from favoring or discriminating against Internet traffic. You've got Commissioner Tate's working sort of definition, which wants to add to the AT&T merger a prohibition on AT&T charging websites for delivering their content to Internet users.

As my friend James Gatusso at the Heritage Foundation has pointed out, what's at stake here is whether an Internet provider can prioritize different bits of information. As messages and websites are broken down as they travel through the Internet, can you create a priority and charge a fee for faster delivery? In the end, there are many folks -- Randy May, who I've worked with, mentions this all the time--who wonder if you make Net Neutrality policy, how do you avoid then taking the next step in treating Internet providers as common carriers, and with that the concomitant rate regulation? These are questions our panelists will address and educate us on.

I'm going to read you a brief little bio on each of them, and then let them talk in the order that I'll introduce them. Our first panelist is former Attorney General William Barr. Bill is currently the Executive Vice President and General Counsel of Verizon, where he heads the Legal, Regulatory, and Government Affairs Group. Prior to that, he was with Bell Atlantic and GTE. He has also practiced law in Shaw Pittman and has served at the Justice Department, both in the first Bush administration, ultimately as Attorney General, and then served in the Reagan administration in the White House.

Our second speaker is Paul Misener. Paul is both an engineer, having gotten his engineering degree from Princeton, and a lawyer, graduating from George Mason. He is currently Amazon.com's Vice President for Global Public Policy, and he's responsible for formulating and representing the company's public policy positions worldwide. We are glad that he is here joining us today on this debate. At an earlier point in his life, he was also a practicing attorney at Wiley, Rein and Fielding.

Our final speaker will be Professor Christopher Yoo. Professor Yoo is no stranger to Federalist Society conferences. He is currently a professor of law and Director of the Technology and Entertainment

Law Program at Vanderbilt University, although he informed me that he is now planning to move to the University of Pennsylvania. Vanderbilt's loss will be Pennsylvania's gain. He has done a lot of work in this area, in particular in his book, *Networks in Telecommunications, Economics and the Law*.

So I will now turn it over to the panelists.

William P. Barr: Thank you very much, Dave. It's a pleasure to be on this panel this afternoon with my distinguished co-panelists.

It is no accident that the broadband revolution was launched on deregulated networks. It started in the last decade, with the cable companies investing over $100 billion to convert their broadcast pipes into two-way broadband pipes. For a while, cable dominated the field because phone companies were subject to a regulatory regime that required them to share the lines and to sell their services at regulated rates.

It was precisely when those rules were lifted that phone companies started making substantial investments. Investments soared in broadband. DSL deployment has sharply ramped up since deregulation, and DSL is now gaining share on cable modem's growth. Now, the phone companies are moving to the next generation of broadband. Verizon is spending $18 billion to deploy our fiber-to-the-premises system to 18 million customers by 2010. AT&T is spending $4.6 billion over the next three years to deploy fiber-to-the-node to 19 million homes. Wireless companies are, meanwhile, investing in their wireless networks to deploy 3G broadband technology. Verizon Wireless has spent $3 billion so far to reach 200 million people by the end of this year.

Fixed wireless has now become a viable broadband alternative. WiMAX reportedly will allow speeds up to 155 Mbps over a range of 30 miles. Clearwire, with Intel's backing, is now offering WiMAX in 30 cities and expanding. TowerStream is offering WiMAX in six metropolitan areas. In August, Sprint announced that by the end of 2008 it will spend $3 billion to build a nationwide WiMAX network to provide customers access to the Internet at 2 to 4 Mbps. Several hundred U.S. municipalities are in the process of installing citywide WiFi networks. Already, about 65 municipalities have such networks. The three satellite companies are continuing to invest in substantially improving their nationwide broadband offerings and report that subscribership is increasing. Recent technological advances have now made broadband-over-power-lines ("BPL") a feasible access alternative, and Google-backed Current Communications is rolling out BPL in Texas and Ohio. Current speeds are up to 3 Mbps.

The bottom line is that we have underway probably the largest infrastructural deployment in recent history. Over the last two years, Verizon has been the number one capital spender in the country. Unlike most historic infrastructural projects of this scale, however, these builders are not being granted exclusive franchises and promised relatively safe returns. They're rolling out their networks in fiercely competitive markets, markets that are subject to extraordinary technological risk. When Verizon puts fiber down a street at the cost of about $850 per home passed, we do not know whether any customer on that street is going to sign up for our service. And when we drop a line to the house at roughly $1,000 per home, we have no idea whether that customer is going to turn to cable, WiMAX, or some other competitor shortly thereafter.

Again, it's no accident that these investments are being made in a deregulated environment because companies are going to make these kinds of investments only if they see an opportunity to earn a return that is commensurate with the risk, and only if they have the freedom to innovate, differentiate, and make commercially sensible decisions that they need to compete and win in the market.

Let's consider now the argument of the advocates for broadband regulation. Their basic claim is that the market for last-mile Internet access is really a duopoly controlled by cable and phone companies with enough market power to harm competition in the market for content and applications. Their prescription is a set of *ex ante* blanket rules governing the way business is transacted on the Internet in a comprehensive regulatory scheme.

Now, the claim is frequently made in the public debate that what the last mile providers are going to do is block or interfere with the content at Yahoo or eBay or Amazon or the like that people are reaching over the public Internet. But that's not really what the debate is about. We've made clear that consumers should be able to reach any lawful website that they want with the access service they have bought, and

we do not and will not block, degrade, or interfere with the consumers' access to those websites. Indeed, one of the reasons we're making the investments we are is precisely to enhance the experience and the range of services that our end users can get over the public Internet.

The real issue is this. Broadband pipes not only enhance what people can get over the public Internet. They also make possible new kinds of priority delivery services, quality of service capabilities, and functionalities that by definition cannot be accommodated on the public Internet. The fact is that the public Internet as it is now configured could not seamlessly accommodate the explosion of gaming, on-line movie viewing and other bandwidth-intensive applications that are envisioned for the future. The question is whether we are going to be able to build the infrastructure and to develop these new functionalities and capabilities and offer these new services to businesses, so they can in turn offer to end users new services that otherwise would not be available at all.

So, for example, if Johns Hopkins Hospital wanted to develop and deploy a home monitoring network by which it could monitor very sick patients at home and provide certain medical services remotely right to the patient's home, we should be able to deploy for them a network with very high degrees of quality of service, security, reliability, and end-to-end management of their traffic. The so-called net neutrality advocates say that if a network owner wants to provide a new enhanced service, they should only be able to charge the end user; they should not be able to charge the company that wanted to provide these new services. The result is that the company that wants to provide new services for its customers would be banned from working to help the network company build this new service.

Some say that if the network provider is going to provide new enhanced service to some content or applications providers, and do it for a fee, then the network provider has to provide the identical service to all comers on exactly the same economic terms. Some add that the network owner must be prohibited from providing itself any functionality or capability that it does not make available to all comers on the same terms. These non-discrimination models require, as they always have, intense regulatory oversight of all the physical and economic terms of transactions. Some regulator will have to determine which parties are similarly situated, what kind of businesses are equivalent, what kinds of terms are equivalent, what different portions of deals could be carried over to another deal, etc. But more importantly, non-discrimination regimes like this ultimately lead to regulated prices, that is, to tariffed rates set by regulators.

Obviously, in our system, the presumption is against regulation. The burden is on those seeking regulation to show that in fact there is a market failure causing harm to competition; and, moreover, that the regulation will actually improve things, not make them worse.

The threshold problem with the broadband regulation argument is the harm it posits. Network providers simply lack the market power to harm competition in the content and applications market. No phone company or cable company has the market power to injure competition among content and applications providers.

The suggestion that this is a duopoly is an exaggeration or misrepresentation. The broadband market is fiercely competitive today, and its trajectory is to become even more competitive. As noted earlier, multiple technologies deliver broadband services. And many of these technologies—such as WiFi, satellite, and WiMax—have comparatively modest build-out costs. The result is that barriers to entry are comparatively low, while the incentives for entering the market are high. Thus, consumers have multiple choices of access providers, and the choices are rapidly expanding. Eighty-one percent of zip codes have three or more choices. Fifty-three percent of zip codes have five or more choices. Twenty-one percent have ten or more. Broadband prices clearly do not reflect market power. On the contrary, they have been trending downward very sharply, and speeds have been increasing. DSL prices have fallen nearly 30 percent in three years, and by nearly 50 percent at any given speed. And cable modem prices have decreased 70 percent in three years on a per-Mbps basis.

Moreover, advocates of regulation are engaged in a sleight-of-hand here as to what the relevant market is. The broadband regulation argument hinges on the power of the last-mile provider over the upstream content-and-applications market—and that is a national or global market. Whatever Verizon's share today in a particular city may be, it only has 12

percent nationally and 2 percent globally. Because of this fractured structure of the industry, no last mile provider has power over the national or global market for distribution of content and applications.

Moreover, last mile providers such as the telecommunications and cable companies lack the incentive to limit the experience of their end users on the public Internet. Indeed, we are selling—and actively promoting—the ability to reach Internet websites. In a word, we sell access. That's our primary business. It is in our interest, obviously, to increase the value of what we are selling by maximizing the amount of content and applications that are available to the end user. At the same time, we know that end users will shun a system that restricts their access to desired content. No network can succeed in the broadband marketplace if it acquires a reputation for not delivering on customers' universal expectation of access to the content they want.

It should not surprise you that when given the lack of market power or the incentive, the history is clear: broadband regulation advocates cannot point to any of the harms that they are concerned about as actually having materialized in the marketplace. The one paltry poster child that they wheel out really shows the vacuousness of the claim. In 2005, Madison River Communications, a small rural telephone company with 190,000 lines, blocked Vonage, a VoIP provider, from terminating on their system, apparently over concerns that Vonage was not paying them access charges. This was a legal dispute about a question that is still being contested: does VoIP traffic have to pay terminating access charges as other long-distance companies do? In any event, the FCC staff quickly reached an agreement with Madison River whereby the company paid $15,000 and agreed to stop blocking the calls. That is not a predicate for the kind of massive regulation that is being called for.

In short, the broadband market is characterized by multiple competitors, falling prices, increasing transmission speeds, new investments, and vibrant innovation, all characteristics of a marketplace that is not in need of intervention by regulators. A fundamental problem of these *ex ante* regulations that are being proposed is that, as Professor Chris Yoo has pointed out, they are addressed to the wrong problem. The premise of net neutrality regulation is that our policies have to be targeted to fostering competition in the content-and-applications market. But that market already is highly competitive and becoming more so. If, however, as some regulation advocates suggest, the problem is concentration or scarcity at the network level, then the policy imperative should be to broaden the availability of network capacity and network capabilities by promoting investment in multiple diverse networks. If the problem is too few networks, the solution is more networks.

Even if the ultimate concern is the content-and-applications level of the marketplace, it is still the imperative in the first instance to give priority toward policies geared toward encouraging the deployment of diverse networks. It is hard to imagine a wealth of new applications being written if there are no networks to support them. And each time a network owner invests and innovates to create a new network capacity and function, it enables a whole spectrum of content and applications that did not exist before.

The fallacy of the regulation approach is that it posits the problem of network scarcity but fails to address it. To the contrary, it assumes that enduring network scarcity is a given and prescribes a regulatory scheme that carves up network resources to all comers, either for free or at regulated rates. The problem, of course, is that these very regulations will deter the building of new networks by severely constraining the ability of network owners to innovate, differentiate, and earn a return that justifies investment in networks. This ends up locking scarcity into place and stunting the market.

It is critical to understand that today's network infrastructure will not support the rich array of content and applications that are on the drawing boards. And the problem is not just capacity on the last mile or capacity on the backbone. It goes to the very network functionality of the public Internet. The real constraints to applications, right now, are the limitations that have been built in to the public Internet. The public Internet gives no set of bits priority over any set of bits. It also operates under the regime that the level of service that you can provide is best efforts. It does not allow for quality of service on the Internet, and this could preclude many types of next-generation content and applications.

Now, this does not mean that the public Internet is going to be superseded. On the contrary, it is going to remain the primary delivery vehicle for most of the kinds of consumer content and applications with

which we are familiar —and that is exactly why we are investing in making the backbone more robust and making the last mile robust. But it does mean that if we are going to expand the universe of applications, we are going to have to promote the diversification of networks: networks that allow prioritization, networks that allow a range of quality of service, networks that are optimized for particular kinds of content and applications. And the regulatory regime that is being promoted does the opposite. It deters investments, and it threatens the ability to recoup investment and deprives the network owner of the freedom to compete. The proposed regulatory regime equals a less diverse Internet. It means not more but fewer products, services and players.

Take, for example, the dictate that network providers should not be able to charge upstream providers for new and enhanced services. This is intuitively and obviously wrong. If we do something with Johns Hopkins Hospital, why should only the patients pay? The infrastructure and new functionality we are building makes markets. These markets are two-sided markets. There are times when customers want to pay to reach the businesses, and there are times that the businesses want to pay to reach the customers.

Take the market for express delivery served by companies like Federal Express and UPS. There are times when consumers are willing to pay for express delivery—for example, when I need a book delivered overnight—but there are also times when the sender is willing to pay for express delivery—as, for example, where a law firm has to get a brief to a court immediately. The idea that only customers should be able to pay for delivery forecloses large avenues of efficient activity and eliminates key revenue sources for network providers to recover their investment.

Content and applications providers like online gaming companies may be willing to cover some of the expenses of improved delivery services and to help make a market for their services. Ultimately, the nirvana for gaming is to have virtually no latency in the system, so that when I make a move on one side of the earth, it is almost instantaneously perceived on the other. That requires a very robust network with a very high degree of quality of service and prioritization. A gaming company may want to enter into a transaction with a network company to provide just that kind of capability. For example, perhaps when the user accesses a game and wants to play it, a burst of capacity is made available on that site from the gaming company to allow that to happen. There is no reason why these kinds of arrangements should not be allowed to take place. This improves competition on the network level, and it improves competition at the content-and-applications level.

Some proponents of Internet regulation concede that network providers should be able to charge upstream for these services, but demand that if a network owner provides this for anybody, it has to provide for everybody on the same terms. That kind of requirement is unnecessary and harmful. First of all, network providers have incentives to maximize the diversity of content and applications on their networks, as I have described. Moreover, once it is conceded that a network provider can negotiate a commercial arrangement with a content provider that wants to reach the customer, what is the reason for not allowing those very same market forces to govern the transactions with the second, third, and fourth content provider?

The problem with non-discrimination requirements is the certainty of regulatory failure. Enforcing non-discrimination obligations requires the regulator to determine which providers are similarly situated. Take for example the Johns Hopkins hypothetical I described earlier, in which Johns Hopkins traffic was prioritized from end to end (including the last mile). Suppose another company —say, Victoria's Secret—notes that Johns Hopkins traffic was given this prioritization in the last mile and demands the same last-mile prioritization. That requires the regulator to price that last-mile bit of prioritization, which was but one component of the whole Johns Hopkins transaction. Isolating that one aspect of the larger deal, defining it, and pricing it are devilishly hard to do, making the risk of regulatory failure here alarmingly high.

Moreover, if the advocates for broadband regulation believe that this is not a competitive market, that we need these rates and regulators because there is a bottleneck, are they really going to be satisfied with a commercial, real-market rate? Or, as in every case in telecommunications until now, will they require the regulator to determine what the rate should be? How this is going to be done in a competitive market is beyond me. The regulators have shown that they cannot do this. This will end up in

undercompensation of the network companies. The bottom line is that network companies are unlikely to take the kinds of risks and make the massive kinds of investment going on today if all they can expect are regulated rates for their return.

If a problem arises, there are rules in place sufficient to deal with it. The FCC has said that it has authority over this market, and it stands ready to come in if it sees any abuses or anticompetitive conduct. That's the way it should be. Until actual harm arises that can be reviewed in context to see whether it harms competition, whether there is market power, etc., there is no reason to impose *ex ante* a broad set of rules that locks in concrete how the Internet—and commercial relationships on the Internet—should proceed.

Paul Misener: Before I begin, thank you very much, Bill. That was a great introduction to your views on this debate. You have also given me an opportunity to speak a lot slower than I had planned. I'm from Amazon.com, and my remarks today really do represent only Amazon's views on this matter. There are other companies we've worked with over time that may not agree with everything here, and I hope there's enough substance to show that we're a little bit closer than perhaps is imagined.

So let's quick-start with how the Internet works. I'll talk about what net neutrality means to us and really describe the meat of this issue, how it is a disagreement on the facts, not so much on the slogans or the philosophy. And then, provide some more discussion of how the net ops get paid; (hopefully, this is a place where Bill and I perhaps can narrow the chasm).

The engineering is really at the core of the policy. The home user has a subscription with a broadband network provider like Verizon. Verizon then is interconnected to the Internet backbone. On the other end is a service provider, a content provider online, like Amazon, Yahoo, or YouTube; they're connected through agreements with a business ISP. How does the content that is provided or made available by the service provider brought to the home user? Well, it's very simple. There's something called the hypertext transfer protocol, which governs how communications are made on the web, and there's a command called "get." So, when you type in a URL on the top line of your browser, or when you click on a link on a webpage, it actually sends what is called the get command. The get command is destined to a particular server. In fact, it's destined to a particular file on that server. And so, when the page comes up and you see it, all that is, is a file. And you go on and you send a get command to have that file sent to you. It's actually called a resource. You've certainly seen this before; the uniform resource locator, the URL. That's the resource. It really means just content, be it the webpage, be it a video, an image, text, whatever. But the important policy point here is that the resource, that is the content, only gets into the Verizon or Comcast network if the home user who has paid for the access has asked for it. This is unlike every other medium in history, where it was decided by the publisher, the broadcaster, the writer, the author, what gets sent out to the consumers. Here, the consumer chooses. If no consumer ever chooses the particular content on the other side of the web, that content never gets into the Verizon network. Interesting, huh?

So, how do net ops get paid? Well, this is really basic stuff. Right now, they are paid. There's a network operator on the consumer side who gets paid by the consumer. It is not a fixed rate necessarily. It's been priced that way in some markets, but you can see the differences in prices between DSL and cable modem access, which is explainable largely by the different speeds that they provide. DSL is cheaper because it's slower in those markets. Likewise, the other ISP, which may or may not be one of the major residential ISPs, gets paid by the business. Amazon pays a lot more for access than Joe's bookstore.com.

The residential broadband network operators want to introduce a second way to charge, to sort of reach through the web and be able to get money for capacity, for content transiting their networks. And they want to charge the source of that content, even though the source of that content didn't put it into the network; it was just made available and pulled by the user who's already paying for it. That's not the end of the diagrams, but we'll get back to them in a second.

It's funny; I've heard so many times how no one knows what net neutrality means, everybody disagrees on it. You know, that must mean it's amorphous and can't be regulated or legislated. Well, that's true of everything. What is the war on terror? What is health care? What is any number of the much

more complicated things that Congress and the FCC deal with? Net neutrality is actually relatively simple compared to the big issues the Senate and House hit everyday.

The core idea of net neutrality is nondiscrimination based on the source or ownership of the content. I've underlined those two words because, you'll notice, it doesn't say the kind of content or the timing of content or the particular technical needs of the content, but rather, the source of the ownership. So, if AT&T wanted to prioritize all Internet video over all Internet data files, fine. That makes perfect sense. That's a rational network management decision. But to choose among the providers puts the network operator in the position of deciding for consumers which content gets favored over other content.

I've heard the Johns Hopkins example before, and I'm a big fan of that particular institution. But the fact of the matter is, you don't want to have a circumstance where the network operator cuts a special deal with Johns Hopkins that it forbids Mayo Clinic. Otherwise, Verizon steps into the position of being an HMO for all of its customers. You don't want that. You don't want Verizon HMO. And, in that sense, net neutrality is really about preserving the openness of the Internet, which has been so great for consumers and innovation.

I want to draw a distinction here, because none of us—well, probably few of us in this room—are great fans of the '96 Act. I share a lot of sympathy with Bill on that particular point. But the '96 Act was about busting up market power. It was trying to dismantle the market power of the telephone companies. This is not what we're seeking to do. We acknowledge that they have market power. We say its okay for them to have market power, so long as that power over the network is not extended to market power over content in ways that have never been done before. We're not seeking to bust them up, just to prevent them from extending the power that's extant.

So, the disagreements is just kind of a bunch of slogans. Believe me, my side of the debate has been guilty of sloganeering also. I get that, and I'm trying to distinguish Amazon's particular viewpoints from everything else by being very specific about what we mean. Let's walk through a couple of these. First, "Content shouldn't fill the pipes for free." Perhaps some of you are familiar with some of the comments made by the leadership of network operators who say they don't want their pipes being used by companies is like Amazon for free. Well, we agree.

The content, as we learned on the very first couple of slides, injures the pipe only when the paying customers go and get it. We're not pushing it out there. We're not a cable content provider that is pushing it out to that set-top box and filling up their pipes, the network operators' pipes. The only reason our content gets there is when they're paying customers ask for it.

Next: "Let the competitive free market work." Well, I am a free market guy, too. I agree: let the competitive free market work. The problem is: there isn't competition. With all due respect to Bill, it is not a competitive market at all. Over 98 percent of residential broadband access is provided either by the phone company or the cable company. These other nascent technologies he discussed are interesting, and they're going to be great at some point hopefully, but they're nowhere being relevant players in the market.

Bill talked a little bit about the zip codes. Well, the fact of the matter is that people don't buy their broadband Internet access for a zip code. They buy it for a house. And while there may be many, many providers within that particular zip code, the individual house, the individual consumer, only has either the cable pipe or the phone pipe—at least 98-plus percent of them if you believe the FCC.

What are some of the other slogans we hear a lot? "We shouldn't start regulating the Internet" Again, I agree. I'm for that. But the fact of the matter is that nondiscrimination rules govern most of consumer Internet access—in fact, by far most and arguably all, because it was not clear for a while. When the Commission started to look at reclassifying broadband access, by far the vast majority of Internet access was under these nondiscrimination rules. So these are historicals. This is a new thing. This would be largely a reinstatement. We can debate about whether it actually applied to cable, but that's not important because the vast majority of consumer access was dial-up at the time.

"Network investments are good for consumers" I totally agree, but that's not to say that the network operators would suffer under a nondiscrimination rule. In fact, in the year preceding the Commission's decision to reclassify (that is, to deregulate) broadband

Internet access, the network operators applied 60 percent more lines in just that one year. They were investing heavily even though the nondiscrimination rules applied. In fact, if you subtract cable from that equation, if you just talk about the ones that were undeniably regulated (that is, the telco providers), it's well over 200 percent more lines in the year before the Commission acted to deregulate. They don't need to discriminate to invest.

"Video competition is good for consumers" An important thing for consumers. But how ironic would it be if, for the sake of getting one more video provider, we cut off access to hundreds or thousands of other ones? I have to say that the net neutrality legislation is really part of a much bigger telecom reform bill that you're probably aware of. This is the original net neutrality Snowe-Dorgan Bill that we put in. This would survive by itself as legislation. This, even thinner bill -- essentially two pages -- was the amendment proposed to debate the Telecom Act. It would accomplish what we want in net neutrality.

This, on the other hand, is the telecom act that Bill's company supports. This is the light regulation, light regulatory touch, the light legislation that he's in favor of. But it's this very heavy legislation that he fears.

Interestingly, there's another slogan that's out there, which says, "Common carriers are bad for consumers," as if somehow net neutrality equated to common carriers. That's simply not true. Nondiscrimination is not all common carriers. There's a lot of bad stuff that used to apply to Verizon that shouldn't apply anymore. But this huge bill, interestingly enough, in many places relies on nondiscrimination rules. They want non-discrimination in law so long as it runs in their favor.

Now, network operators should be paid for their service. This makes perfect sense economically, out of fairness, and for getting a bigger, stronger Internet out there. So, let's talk about this in a little more detail. Basically, you have these content providers (call them Yahoo and Google, or Amazon and Google, whoever you want it to be) to a neighborhood. And the network operator, the broadband residential Internet access provider, be it Verizon or Comcast or AT&T, have all these functional elements to their network. They could be servers. They could be cards. They could even be cached—that is to say in memory -- within a server card. Those are functional elements.

When User A, in his or her home, gets some content from OSP #1, it should in no way interfere with user B's ability to get content from OSP #2. If User A wants high speed or some extra service provided by OSP #1, that's fine. OSP #1 ought to be able to pay for that so long as it never hurts though OSP #2's ability to serve the second home user. That's the basic model. This is how it exists today. By the way, the home user is paying. Both of the users are paying Verizon or Comcast, and then the online service providers are paying for Internet access at their end.

There's something that's done today commonly called edge serving. I don't know if you've ever gone to CNN, but if you watch as the page is loading, you'll see up there in the URL line something that comes up that looks like Akamai. Akamai is a company that provides Web servers at the edge of the network. The reasoning is that, by distributing content around the country in the high population areas, CNN servers in Atlanta don't get hit every time somebody in New York, Detroit, or Los Angeles looks for the latest news on the home page. Most likely it's going to hit an Akamai server that's located in or right around those cities. This goes on today, and if Verizon wants to get into this business, more power to them. That's great. They can do this. This is an example of how the OSPs can pay the residential network operator, like Verizon, more money for enhanced services. But you'll note what it doesn't do. It doesn't in any way interfere with the ability of home User B to get stuff from OSP #1. When it gets to that router device, there's no discrimination.

Here's another thing that happens today. The OSP can sell to Amazon or Google or whomever a private line that skirts the bulk of the networking. It can skirt the cloud entirely, the local network. This is another example of how Verizon can and does sell services to the content providers, who then pay for enhanced service because it doesn't have to go through all these little bumps along the way. But again, when it gets to that router device, there's no discrimination.

A lot has been said about quality of service and how advocates of net neutrality say that we should not allow the network operators to provide quality of service based on the source or ownership of the content. I'm one of those funny guys who actually

disagree. I think it would be okay for quality of service to be sold by the network operators. But there's a catch, and I'll describe this. Here's how it works. This private network here doesn't interfere with the other traffic around there. How would you do that with quality of service? When Bill talked about prioritization, he said we're not going to block anything but we're going to prioritize some stuff; we want to be paid to prioritize. Well, you can't pay for useful prioritization unless everybody else suffers degradation, right? If all the seats in an airplane were first-class, no one would pay extra for them, right? So the very fact of the matter is that you pay to get priority, to get better than everybody else. Otherwise, no one would pay for it. Why would you? This is OK; except that's not how it works here.

How it works here is you're actually going around the network and you're getting better service, but not at the expense of these guys. It's not actually hurting them. And the only reason is because that's a new capacity. It didn't exist before. So you're not subtracting away from OSP #2; you're actually adding something in. So the same thing can be done for quality of service within the network.

Here's how. Why doesn't Verizon offer something new inside the network? Remember that little box; that could be a new router, it could be a new card, it could be better software. What it does is provide priority, a faster service, better speed for OSP #1, the red guy, and it gets through faster to home User A down there. But what it doesn't do is in any way affect OSP #2 and home User B, because it's new capacity. So quality of service paid for by the OSP, by Amazon, by Google, a new capacity within the network, seems fine to me. Why shouldn't it be? It's just a private network.

Likewise, if they wanted to do something called "the turbo button." BellSouth has experimented a little bit with this; the concept is that the home user pays a little bit more to get a boost in speed. That's fine so long as it's at a new capacity; but if it's not new capacity, every time the guy down hits turbo, User B's content gets screwed up, its slowed down. And that's not fair. So, the concept here is—(and I'm positing this as possibly a middle ground for discussion)—is that quality of service and new capacity ought to be acceptable under net neutrality rules. But if it's in the existing capacity, where it hurts other consumers' ability to get at other content on the Web, that's a problem.

I'm going to spend just two minutes, if I may, on answering a couple of Bill's points. First, he said that falling prices are an indication of the competition. Well, no, not necessarily. Firms try to price at the profit-maximizing point, and if they've priced too high, they can still come down to a profit maximization point, or closer to it—(at least firms with market power can)—and still drop their prices. So it's actually been suggested that that is probably what's going on here. What is it, something about 85 percent or so of homes are passed by residential broadband Internet access? And about 40 percent —(that's probably forgiving)—take it? That's a huge gap of people who could get it but don't. Why? The vast majority say it's too expensive. Well, I think right now prices are falling in part because they view the ability to pick up more consumers, even if their subscription profits from existing consumers decrease slightly. It gets them, again, to the profit maximizing point.

A lot of them inaudible that DSL is cheaper than cable. Well, in many respects it's a different service, as I mentioned before. It's slower, so you want to pay less for it. That makes perfect sense. That's buying by quantity. But this business about applauding this slight drop in prices for broadband Internet access as being evidence of competition strikes me as more or less like the policeman who's pulled you over—doing 65 in a 25-miles-an-hour zone. You say to the policeman, "Well, last week I was doing 75; you should be happy." The fact is, we're not even near the competitive price, and it's because we've got this very powerful duopoly.

Again, I do believe it is based on facts, the disagreement between us. And I think if we really look hard and decide what the facts actually are—(and I'm trying to base my views on published sources like the FCC)—we'll get closer to a solution, because, as Bill acknowledged, much of our argument, the pro-net neutrality folks' argument, rests on competition. We at Amazon believe that when there is a demonstrable level of competition, some sufficient level—and we can argue about what that is—that would be the end of such rules. There would be no need for the regulation when a truly competitive market is in place. I think there are others perhaps in Washington who believe this should last in perpetuity. I don't. Amazon doesn't. And so I hope we can, at some point, sit down and agree on the facts and possibilities

like this way of providing quality of service and new capacity, and then hopefully get beyond just the slogans.

Thanks very much.

Christopher Yoo: In many ways, the debate over network neutrality is the direct result of the dramatic changes in the nature of the Internet over the last decade. The current Internet is effectively standardized on a suite of protocols known as TCP/IP, which for purposes of the network neutrality debate has two distinctive features. First, it routes traffic on a "best efforts" basis without any guarantee that any particular packet will ever arrive. Second, it routes traffic on a "first come, first served" basis that does not give priority to packets associated with particular content or applications.

This approach worked fairly well when the Internet was primarily a means for academics to exchange e-mail and text files, in which delays of less than a second were virtually unnoticeable. Starting in the mid-1990s, the privatization of the NSF backbone and the accompanying elimination of the commercialization restrictions transformed the Internet into a mass-market phenomenon. The number of people using the Internet exploded, which in turn caused an exponential increase in the number of possible connections. The emergence of new applications also caused a dramatic increase in the heterogeneity of network usage. It is only natural that the Internet would evolve to meet these new demands. Consider, for example, Internet telephony (also known as voice over Internet protocol or "VoIP"). The International Telecommunication Union standard requires service with latency of no more than 0.3 seconds. Anything less renders telephone service unusable.

Furthermore, graphics-intensive applications, such as video and graphics-intensive online gaming, require more bandwidth than e-mail and web browsing and are often exhibit greater variability of demand, which in turn makes it all the more important to permit network owners to experiment with new approaches to network management. One solution would be to increase bandwidth. Another solution would be to give a higher priority to the traffic associated with applications that are sensitive to delay. Still another solution would be to cache content at multiple locations around the Internet, as is currently done by content distribution networks like Akamai. Which solution will represent the most efficient approach at any time will depend on their relative costs. The law of diminishing marginal returns dictates that the marginal gains from any one approach will eventually tail off to the point where some alternative architectural solution becomes preferable. Technological change can also cause costs to change in ways that may change the costs and benefits associated with any one approach. There thus seems no reason to presume a priori that any one approach will emerge as the best solution in every situation.

Mandatory access requirements like network neutrality threaten to limit network owners' ability employ alternative approaches to network management. As I have noted in my earlier work, access requirements entail the adoption of four corollaries. First, regulators must require the network owner to permit third parties to interconnect with their networks. Second, the regulatory scheme must define and standardize the interface through which interconnection must occur. Third, because the network owner could render any access requirement a dead letter simply by charging unaffiliated content and applications providers more for access than it charges to its own proprietary services, any access regime must also include a nondiscrimination requirement. Fourth, access requirements necessarily entail some form of rate regulation. This is because a network owner could charge a nondiscriminatory price and still effectively exclude unaffiliated content and application providers simply by charging uniformly exorbitant prices. Such a price would have no real impact on the network owner's bottom line, as it would simply transfer profits from the content and applications subsidiary to the last-mile subsidiary.

Network neutrality would thus necessarily require the imposition of a fairly intrusive regulatory regime that includes elements that have proven extremely difficult to implement in the past. Furthermore, the standardization and interconnection requirements threaten to retard innovation by locking the existing interfaces into place. Under the best of circumstances, the requirements of the Administrative Procedure Act dictate that any adjustments to the interface would take a minimum of several months. The need to preserve such experimentation is what has led a growing number of senior network

engineers, including TCP/IP co-author Robert Kahn, end-to-end co-author David Clark, and the so-called "grandfather of the Internet" David Farber, to oppose network neutrality. Even worse, mandating access through regulation would politicize the decisionmaking and would render access vulnerable to the well-recognized defects of the administrative process revealed by public choice analysis.

In addition to expecting network owners to employ a broader range of techniques for managing network traffic, the increasing heterogeneity of network usage should cause pricing to become more complex. To date, pricing on the Internet has been relatively simple. Networks have traditionally offered end users "all-you-can-eat" pricing that charges a single, flat fee that does not vary with the amount of bandwidth consumed. Economically rational end users will increase their network usage until the marginal benefit they would derive from any further increases no longer exceeds the marginal cost of doing so. Because under all-you-can-eat pricing the marginal cost of increasing consumption is always zero, end users continue to increase their consumption so long as they derive any positive benefit, no matter how small. The problem from a social welfare standpoint is that increases in consumption impose congestion costs on other users that are not taken into account when individual users calibrate their demand. The wedge between private cost and social cost gives end users a systematic incentive to overconsume.

One logical way to eliminate this problem is to charge end users a usage-sensitive price set equal to their marginal contribution to congestion. The problem is complicated by the fact that determining the congestion costs created by a particular user at any particular time can be quite complex. As an initial matter, if the relevant portion of the network is slack, the marginal contribution to congestion may be essentially zero. The situation is more ambiguous if the relevant portion of the network is saturated. The network may be able to accommodate additional traffic by rerouting it along different pathways, depending on the other traffic in the network. If large portions of the network are close to saturation, it is also quite possible that the increase in congestion will cause a cascade effect that amplifies the impact of the increase in congestion. Rationalizing consumer behavior by charging the usage-sensitive fees precisely calibrated to their contribution to congestion would thus require a dynamic pricing scheme that varied depending on the particular configuration as well as the volume and pattern of other traffic passing through the network.

Network neutrality is not only questionable from the standpoint of network management; it may well harm consumers. A network owner that charges all end users a single, flat price would naturally set that price equal to the cost imposed by the average user. This would effectively require low-volume users, who impose less than average congestion costs on the network, to cross subsidize the high-volume users, who only pay the average contribution to congestion even though they are responsible for a disproportionate amount of the congestion.

There is also no reason to expect that the pressure to make pricing more complex will be limited to the side of the market in which last-mile providers bargain with end users. The key to this insight is recognizing that last-mile providers operate in a two-sided market. On one side of the market, they bargain with end users. On the other side of the market, they bargain with content and applications providers. We should expect pricing with respect to content and applications providers to become more complex as well. Some content providers, such as bloggers, primarily transmit text, which uses relatively little bandwidth and is not particularly sensitive to delay. Other content providers, such as providers of streaming video, have much higher bandwidth requirements and require guaranteed levels of quality of service. If network owners are not permitted to experiment with differential pricing, they will be forced to fund any network improvements by charging a uniform price to both types of users, even though the non-bandwidth-intensive providers do not need and do not use the additional network capabilities. The far more sensible approach would be to permit network owners to charge more to those content and application providers that benefit from the network improvements without having also to charge more to those providers who were perfectly happy with the network's capabilities before it was upgraded.

Furthermore, the fact that this is a two-sided market means that the prices charged to end users and the prices charged to content and applications providers are linked in a fundamental way. Consider what would happen if, as some network neutrality

proponents suggest, network owners were prevented from varying the amount they charge particular content providers and application providers, but were permitted to vary the amount they charge particular consumers. Such a regime would limit network owners' ability to extract surplus from content and applications providers, while giving them greater ability to extract surplus from consumers. The net effect of such a regime would have the somewhat perverse effect of forcing consumers to bear a greater proportion of the costs of network improvements, such as building fiber to the curb.

The most serious problem, however, is that the network neutrality debate is focusing on the wrong policy problem. According to the economic theory, any vertical chain of production will only be efficient if each level of the chain of production is competitive. This in turn suggests that the central focus of competition policy should be to identify the level of production that is the most concentrated and the most protected by entry barriers and attempt to render that level more competitive. In the case of the Internet, the level of production that is the most concentrated and protected by entry barriers is almost certainly the last mile. That being the case, one would expect the network neutrality debate to turn on how best to promote competition in that segment of the industry. Instead, network neutrality proponents direct their proposals on how to maintain and promote competition in content and applications, the segment of the industry that is already the most competitive, the least protected by barriers to entry, and thus the most likely to stay that way.

The proper focus of the debate should thus be on the impact that mandating network neutrality would have on the competitiveness of the last mile. In the past, access requirements (such as the unbundled access requirements established by the Telecommunications Act of 1996 and the "equal access" mandate created during the breakup of AT&T) were imposed when competition in the last-mile was believed to be infeasible. As a result, policymakers and courts abandoned the first-best goal of promoting competition in the last mile and instead pursued the second-best goal of promoting competition in complementary services, such as long distance and information services. The infeasibility of last-mile competition rendered the fact that access requirements deterred investment in alternative last-mile technologies of little import.

The analysis changes dramatically once competition among alternative last-mile providers becomes economically viable. Once that occurs, the proper course of action is to return to the first-best policy goal of promoting competition in the last mile. Commentators have long recognized how access requirements can dampen incentives to invest in alternative network technologies, by rescuing those denied access to the existing network from having to invest in alternative sources of supply, which in turn deprives those seeking to build those alternative networks of their natural strategic partners. This dynamic is eloquently demonstrated by the conduct of device manufacturers and content and applications providers after the Supreme Court's *Brand X* decision made clear that FCC regulation would no longer guarantee access to existing last-mile broadband networks. Immediately after *Brand X* was decided, Disney, IBM, Intel, and others began pouring money into new last-mile technologies, such as broadband over powerline and wireless Internet. Most dramatically, Google promised to build a wireless broadband network for San Francisco for free. This was not an act of corporate charity. Faced with the alternative of being cut off from the network that exists today, these companies began investing in creating the network of tomorrow.

Vertical integration theory also suggests that network neutrality is unlikely to yield consumer benefits in terms of price. According to standard oligopoly theory, the prices charged on each side of the two-sided market depend on the relative bargaining power of the parties, which in turn depends on the number of available alternatives. Prohibiting network owners from discriminating in the upstream market in which they meet content and applications providers will not alter the number of options available in the market in which they meet end users. More concretely, in my house in Nashville, Tennessee, I essentially have only two last-mile broadband options: cable modem and DSL. I would still have the same number of options even if network neutrality were imposed. As a result, I would not expect the prices charged by last-mile providers to change one whit. Imposing network neutrality would, however, have a dramatic impact on the bargaining power in the upstream market in which network owners bargain with content and

applications providers, in which they determine how they divide up the rents extracted from end users. Although the division of those rents are of acute interest to the shareholders of those companies, it is not ultimately a policy problem. In that sense, network neutrality is less about protecting consumers and is more a battle between the Comcasts and the Googles of the world.

Lastly, allowing different networks to pursue different networking strategies allows them to compete on dimensions other than price and network size, which are considerations that favor the largest players. Increasing the number of ways in which networks can compete with one another can make it easier for multiple networks to survive notwithstanding the scale economies created by large sunk costs and network economic effects. I can see a world in which three last-mile networks can coexist: one optimized for current applications such as web browsing and e-mail; another using priority-based routing to facilitate delay-sensitive applications like VoIP; and a third focused on providing security to facilitate e-commerce. Allowing this type of network diversity allows smaller networks to survive in much the same way that specialty stores survive in a Wal-Mart world. by targeting subsegments of the market that place a particularly high value on a particular type of network service.

There would thus seem to be good reason not to erect categorical restrictions that would prevent network owners from experimenting with alternative pricing regimes and alternative approaches to network management. To say that deviations from network neutrality can be economically beneficial is not to say that they will necessarily be beneficial in every case. Although modern economic theory indicates that integration of content and conduit will rarely harm competition, the post-Chicago literature has identified the existence of narrow circumstances under which vertical integration can harm competition. It is for this reason that the literature and the doctrine has never embraced calls to treat vertical integration as legal per se.

Fortunately, the Supreme Court's antitrust jurisprudence provides a useful guidance on how to proceed. Under this approach, practices are categorically prohibited only if they evince such a "pernicious effect on competition" and such a "lack of any redeeming virtue" that nothing would be lost declaring them illegal without requiring any demonstrable harm or inquiring whether any efficiencies exist that might justify the practice. When particular practices may be either economical beneficial or detrimental, the Court has refused to prohibit them categorically. Instead, it has permitted those practices to go forward until concrete harm to competition can be demonstrated in a particular case. Barring a practice only after a concrete harm to competition has been demonstrated gives technological and economic progress the breathing room they need to move forward.

The network diversity approach that I am advocating would thus forego *ex ante* regulation in favor of an *ex post* case-by-case approach in which the burden of proof rests on those challenging the practice. In particular, my approach would require proof of concentration and barriers to entry in the relevant markets. It would also require the articulation of a coherent theory explaining why a particular network owner has the incentive to discriminate against particular content and applications providers in a manner that harms competition. For example, network owners that do not offer their own auction sites have no incentive to discriminate against eBay. On the contrary, they can be expected to embrace eBay as the best method for maximizing the value of their network to their subscribers. Conversely, a telephone company may have some incentive to discriminate against VoIP; but would have no incentive to discriminate against services that they do not offer, like streaming video. Cable operators may similarly have incentive to disfavor alternative sources of video content, but have no incentive to reject technologies that allow them to provide voice service.

In short, even if the concerns raised by network neutrality proponents are taken to heart, they would not support imposition of a general network neutrality rule requiring network owners to provide nondiscriminatory treatment for all content and applications. At most, they would support a targeted rule limited to content and applications that competes directly with proprietary services offered by the network owner. Any expansion beyond that scope would impose regulation even in the absence of a coherent theory of why the market is likely to fail.