Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like

Barbara van Schewick
Stanford Law School

Public Law and Legal Theory Working Paper Series
Research Paper No. 2459568

John M. Olin Program in Law and Economics Working Paper Series
Paper No. 462

Forthcoming, Stanford Law Review, Volume 67, Issue 1 (2015)

Stanford Law School
Crown Quadrangle
Stanford, California 94305

Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like[1]

Barbara van Schewick

June 26, 2014

Forthcoming, Stanford Law Review, Volume 67, Issue 1 (2015)

[1] Barbara van Schewick is a Professor of Law, Helen L. Crocker Faculty Scholar and Director of the Center for Internet and Society at Stanford Law School and a Professor (by Courtesy) of Electrical Engineering at Stanford University. This paper is a revised version of a paper that was published as a White Paper by Stanford Law School's Center for Internet and Society in June 2012 (van Schewick (2012b)). Earlier versions of this paper were filed with the Federal Communications Commission in August 2010 (van Schewick (2010e)), presented at the 38th Research Conference on Communication, Information and Internet Policy (TPRC 2010) (van Schewick (2010g)) and filed with the Federal Communications Commission in December 2010 (van Schewick (2010f)).

Many thanks to Marvin Ammori, Scott Bradner, Ryan Calo, Richard Clarke, Alissa Cooper, Mariano-Florentino Cuéllar, Antony Falzone, Laura Fong, Brett Frischmann, Martin Geddes, Jennifer Granick, Mark Kelman, Dan Kessler, Mark Lemley, Peter Menell, Luke Pelikan, David Reed, Morgan Weiland, and to participants of the Stanford Law School Faculty Workshop, the 2010 Wharton Colloquium on Media and Communications Law and TPRC 2010 for their comments on earlier versions of this paper.

Over the past years, I have had many conversations and discussions which have considerably influenced my thinking on these issues. For discussions of the ideas in this paper, I thank Marcelo Bagnulo Braun, Jack Balkin, Yochai Benkler, Brad Burnham, Bob Briscoe, Don Bowman, Gonzalo Camrillo, Vint Cerf, Cindy Cohn, Michael Copps, David Clark, Susan Crawford, Dave Crocker, Leslie Daigle, Neil Davies, Alan Ford, Fred Goldstein, Philip Eardley, Lars Eggert, David Farber, Gerald Faulhaber, Patrick Faltstrom, Ed Felten, Matt Ford, Brett Frischmann, John Giusti, Jeffrey Glueck, Vivek Gupta, Mark Handley, Dietmar Harhoff, Dale Hatfield, Bernd Holznagel, Zachary Katz, Joel Kelsey, Timothy Lee, Benjamin Lennett, Lawrence Lessig, Jason Livingood, Margaret McCarthy, Andrew McLaughlin, Milo Medin, John Mitchell, David R. Oran, Jon Peha, Arnold Picot, Chris Riley, Ben Scott, Aparna Sridhar, Douglas Sicker, Venky Srinivasan, Gigi Sohn, Alan Sykes, David Tannenbaum, Brad Templeton, Steffen van Schewick, Henning Schulzrinne, Richard Shockey, Dow Street, Sally Wentworth, Kevin Werbach, Philip Weiser, Richard Whitt, Matt Wood, Yiannis Yiakoumis, Christopher Yoo, staff at the FCC Chairman's office and in various bureaus of the FCC, the members of the German Commission of Experts on Research and Innovation, the members of the 2009 Internet Architecture Board, the members of the Trilogy Project and the participants of the Münchener Kreis Workshop on Network Neutrality, the 35th Annual IEEE Conference on Local Computer Networks and Workshops, the Annual Business Network and Board of Trustees Symposium at the Santa Fe Institute, the Technical Plenary of the 75th Internet Engineering Task Force Meeting, the speaker series at the Berkeley Center for Law and Technology, Berkman Center for Internet and Society, NYU Institute for Policy Integrity, Princeton Center for Information Technology Policy, Silicon Flatirons Boulder, Stanford Center for Internet and Society, Wharton Legal Studies and Business Ethics Department and Yale Information Society Project.

Many thanks to the librarians at Stanford Law School for their incredible work and to Elaine Adolfo for everything.

ABSTRACT

In December 2010, the Federal Communications Commission (FCC) adopted the Open Internet Order, which enacted binding network neutrality rules for the first time. Network neutrality rules limit the ability of Internet service providers to interfere with the applications, content and services on their networks; they allow users to decide how they want to use the Internet without interference from Internet service providers. In January of this year, the Court of Appeals for the D.C. Circuit struck down the core provisions of the Open Internet Order – the rules against blocking and discrimination. The Court upheld the Open Internet Order's disclosure rule, so Internet service providers still have to publicly disclose any blocking or discrimination that occurs. As a result of this ruling, Internet service providers like Verizon, AT&T or Cox Cable that connect users to the Internet are now free to block any content, service or application they want. They can slow down selected applications, speed up others, or ask application or content providers like Netflix or Spotify to pay fees to reach their users. These practices would fundamentally change how each of us experiences the Internet.

In the wake of the D.C. Circuit's decision, US policy makers have to decide (again), which, if any, network neutrality rules the US should adopt. They essentially have three options: The FCC can take steps to preserve the Open Internet Rules; it can develop a different, narrower network neutrality regime under Section 706 of the Telecommunications Act within the boundaries established by the Court of Appeal's decision; or Congress or the FCC can adopt a new network neutrality regime under Title II of the Telecommunications Act.

This paper helps policy makers think through the available options, focusing on the substantive merits of the different non-discrimination rules under consideration. In addition to rules that forbid network providers from blocking applications, content and services, rules that forbid discrimination are a key component of any network neutrality regime. Non-discrimination rules apply to any form of differential treatment that falls short of blocking. They determine, for example, whether network providers are allowed to provide low-delay service only to their own streaming video application, but not to competing video applications; whether network providers can count only traffic from unaffiliated video applications, but not their own Internet video applications, towards users' monthly bandwidth cap; or whether network providers can charge different Internet access charges depending on the application used, independent of the amount of traffic created by the application.

The paper makes five contributions: *First,* it proposes a substantive framework that policy makers can use to evaluate alternative proposals for network neutrality rules and assess specific forms of discriminatory conduct. *Second*, the paper evaluates eight existing proposals for non-discrimination rules and the Open Internet Order's non-discrimination rule against this framework and proposes a non-discrimination rule that policy makers should adopt around the world – a rule that the Open Internet Order adopted in part. *Third*, the paper highlights the differences between an antitrust framework and the broader theoretical framework on which most calls for network neutrality regulation are based and explains why an antitrust framework does not capture all instances of blocking or discrimination that concern network neutrality

proponents. *Fourth*, the paper offers the first in-depth analysis of the relationship between network neutrality and new network-level services called Quality of Service. *Finally*, the paper provides the first detailed analysis of the Open Internet Order's non-discrimination rule for fixed broadband Internet access.

The paper is relevant beyond the US: In Europe, the European Commission, the European Parliament and the member states are currently considering which approach to network neutrality they should take. The Brazilian Parliament is in the process of adopting network neutrality rules. In all of these debates, non-discrimination rules are a key point of contention. And no matter which network neutrality regime a country adopts, the question of which, if any, network-discriminations require a legal response will remain relevant for years to come.

TABLE OF CONTENTS

NETWORK NEUTRALITY AND QUALITY OF SERVICE: GUIDE TO DISCUSSIONS

LIST OF BOXES

LIST OF TABLES

INTRODUCTION

Who should decide how we can use the Internet? Internet service providers like AT&T, Comcast, Deutsche Telekom, or Telefónica that provide the on-ramps to the Internet, or Internet users? This question is at the core of the debate over network neutrality. Network neutrality rules limit the ability of Internet service providers to interfere with the applications, content, and services on their networks; they allow users to decide how they want to use the Internet without interference from Internet service providers.[1]

The network neutrality debate was triggered by a change in technology. In the original Internet, the network was application-blind, that is, it could not distinguish between the applications, content and services that were running over the network.[2] As a result, network providers could not control the applications and content on their networks. This allowed users to decide how they wanted to use the network, without interference from network providers. Over the past two decades, technology has become available that enables network providers to identify the applications and content on their networks and control their execution.[3]

Network neutrality proponents argue that Internet service providers have incentives to use this new technology in socially harmful ways. They contend that the existing laws in many countries do not sufficiently constrain providers' ability to do so and that, therefore, new rules – so-called "network neutrality rules" – are needed that restrict Internet service providers' ability to interfere with the applications, content and services on their network.[4] According to network neutrality proponents, users, not network providers, must continue to decide how they want to use the Internet if the Internet is to realize its full economic, social, cultural and political potential.

Over the past ten years, few Internet policy issues have received as much public attention. The Federal Communications Commission's Open Internet Proceeding, started in Fall 2009 to realize President Obama's campaign promise to enact network neutrality rules, received more than 100,000 comments from interested parties, many of them ordinary citizens, and was covered extensively in the media, from the *Wall Street Journal* to the *Daily Show*. All over the world, from the US to Europe to Latin America, policy makers continue to investigate whether they should adopt network neutrality rules and, if so, what the rules should be.

---

[1] Throughout this paper, I use the terms "providers of Internet access service", "Internet service providers", or "network providers" interchangeably, and the term "applications" as a shorthand for "applications, content, services, and uses."

[2] The original architecture of the Internet was based on the layering principle and on the broad version of the end-to-end arguments. The Internet's application-blindness was a consequence of this architectural design. There are two versions of the end-to-end arguments, a narrow version and a broad version, which are often confused in policy debates. van Schewick (2010d), pp. 57-81, 377-79; van Schewick (2004). While both versions shaped the original architecture of the Internet (van Schewick (2010d), pp. 90-103, 110-12, 379-81; van Schewick (2004)), only the broad version, together with the layering principle, is responsible for the application-blindness of the network. van Schewick (2010d), pp. 72-75, 217-18; van Schewick (2004). See also Reed (2010). On the layering principle and its relationship to the architecture of the Internet, see van Schewick (2010d), pp. 46-57, 88-90.

[3] On Deep Packet Inspection (DPI) in general, see Anderson (2007). For a specific example, see Cisco Systems (2005a). On the state of DPI deployment, see Free Press (2010), pp. 141-51.

[4] This paragraph is adopted from van Schewick (2009).

In Europe, the European Commission, the European Parliament and the member states are currently considering which approach to network neutrality they should take.[5] The Brazilian Parliament is in the process of adopting network neutrality rules.[6] In the US, the recent decision by the Court of Appeals for the D.C. Circuit in *Verizon v. FCC* has re-ignited the debate. In December 2010, the Federal Communications Commission (FCC) adopted the Open Internet Order, which enacted binding network neutrality rules for the first time.[7] The rules went into effect in November 2011.[8] In January of this year, *Verizon v. FCC* struck down the core provisions of that Order – the rules against blocking and discrimination.[9] The decision combined two wins for the FCC with one decisive loss. According to the Court, the FCC has authority to regulate providers of broadband Internet access service under Section 706 of the Telecommunications Act of 1996, and the FCC's justification for the Open Internet Order is "reasonable and supported by substantial evidence." Both of these points had been heavily contested by Verizon. The Court found, however, that the no-blocking and non-discrimination rules violate the Communications Act's ban on imposing common carrier obligations on entities like Internet service providers that the FCC has not classified as telecommunications service providers under Title II of the Communications Act. The Court upheld the Open Internet Order's disclosure rule, so Internet service providers still have to publicly disclose any blocking or discrimination that occurs.

As a result of this ruling, Internet service providers like Verizon, AT&T or Time Warner that connect users to the Internet are now free to block any content, service or application they want. They can slow down selected applications, speed up others, or ask application or content providers like Netflix or Spotify to pay fees to reach their users. These practices would fundamentally change how each of us experiences the Internet.

In the wake of the D.C. Circuit's decision, US policy makers have to decide (again), which, if any, network neutrality rules the US should adopt. They essentially have three options. First, the FCC can preserve the Open Internet Rules by reclassifying Internet service as a telecommunications service under Title II of the Communications Act. Second, the FCC can develop a different, narrower network neutrality regime under Section 706 of the Telecommunications Act within the boundaries established by the Court of Appeal's decision. Finally, Congress or the FCC can adopt a new network neutrality regime, but only, in the case of the FCC, after reclassifying Internet service as a telecommunications service.

Whether network neutrality rules should include a non-discrimination rule and, if so, what it should be, is a key point of contention in all of these debates. This paper helps policy makers think through the available options. Non-discrimination rules apply to any form of differential

[5] See, e.g.,Kanter & Scott (2014).

[6] See, e.g., Al Jazeera & Reuters (2014).

[7] Federal Communications Commission (2010c).

[8] Federal Communications Commission (2011a), p. 59223.

[9] Verizon v. FCC, No. 11-1355 (D.C.Cir. 2014), available online at http://www.cadc.uscourts.gov/internet/opinions.nsf/3AF8B4D938CDEEA685257C6000532062/$file/11-1355-1474943.pdf.

treatment that falls short of blocking. They determine, for example, whether network providers are allowed to provide low-delay service only to their own streaming video application, but not to competing video applications; whether network providers can count only traffic from unaffiliated video applications, but not their own Internet video applications, towards users' monthly bandwidth cap; or whether network providers can charge their subscribers different prices for Internet access depending on the application used, independent of the amount of traffic created by the application.

The decision for a specific non-discrimination rule has important implications. In particular, it affects how the core of the network can evolve, how network providers can manage their networks, and whether they can offer new network-level services called Quality of Service (QoS). Different applications have different requirements with respect to reliability, bandwidth or delay.[10] While the original Internet provides a single best-effort service for all packets (that is, the network does its best to deliver data packets, but does not provide any guarantees with respect to delay, bandwidth or losses),[11] a network that provides Quality of Service offers different types of service to different data packets.[12] For example, a particular service may guarantee a minimum bandwidth or maximum delay, or it may give some data packets priority over others without giving absolute guarantees.[13] While many applications function well with best-effort service, some applications may benefit from types of service that are more closely tailored to their needs. Whether network providers are able to offer Quality of Service may therefore have implications for the types of applications that the Internet can support.[14]

Thus, policy makers who consider adopting non-discrimination rules face a serious challenge: how to find a non-discrimination rule that realizes the goals of network neutrality regulation without overly constraining the evolution and operation of the network while keeping the cost of regulation low. Overly restrictive rules may impede the evolution of the Internet's network infrastructure in the face of changing requirements, make it more difficult to manage the

[10] For example, Internet telephony is very sensitive to delay above a certain level, but does not care about occasional packet loss. Users usually do not notice a one-way, mouth-to-ear delay of less than 150ms. A delay of more than 400 ms makes voice calls frustrating or unintelligible. (International Telecommunication Union (2003); Kurose & Ross (2010), p. 601.) Depending on the encoding and loss-concealment mechanisms used, Internet telephony applications can tolerate between 1% and 20% of packet loss. (Kurose & Ross (2010), p. 617.) By contrast, e-mail is very sensitive to packet loss, but does not care about some delay. (See, e.g., Kurose & Ross (2010), pp. 92-94 and p. 95, Figure 2.4.) E-Mail applications rely on a transport layer protocol called the Transmission Control Protocol (TCP) to get reliable data delivery. On the needs of applications more generally, see, e.g., Kurose & Ross (2010), pp. 92-95; Peterson & Davie (2012), pp. 530-37.

[11] Thus, the network operates like the default service offered by the postal service, which does not guarantee when a letter will arrive or whether it will arrive at all. Contrary to the postal service, which lets users choose services other than the default service like two-day shipping, the original Internet provides only best-effort service. Peterson & Davie (2012), pp.206-07.

[12] On Quality of Service, see, e.g., Kurose & Ross (2010),pp. 647-72; Peterson & Davie (2012), pp. 530-57.

[13] For example, of the two Quality of Service architectures that were standardized by the Internet Engineering Task Force, the Intserv architecture provides specific QoS guarantees to particular application sessions, while the Diffserv architecture provides different levels of performance to different classes of traffic without providing specific guarantees. On Intserv, see Kurose & Ross (2010), pp. 669-72. On Diffserv, see Kurose & Ross (2010), pp. 660-65.

[14] For a more detailed discussion of Quality of Service and of the conditions under which it is useful, see Section "Ban All Discrimination" below.

networks over which we access the Internet or deprive us of new applications, content and services that new network-level services may enable. Overly permissive non-discrimination rules will fail to realize the goals of network neutrality regulation and prevent the Internet from realizing its economic, social, cultural, and political potential.

This paper addresses that challenge. It makes five contributions:[15]

First, network neutrality proponents generally agree that network neutrality rules should preserve the Internet's ability to serve as an open, general-purpose infrastructure that provides value to society over time in various economic and non-economic ways. A lot of uncertainty exists, however, on how to move from a high-level commitment to network neutrality to a specific set of rules. The paper proposes a framework that policy makers can use to evaluate alternative proposals for network neutrality rules, interpret existing rules, and assess specific forms of discriminatory conduct. In particular, network neutrality rules need to preserve the factors that have allowed the Internet to foster application innovation and economic growth, improve democratic discourse, facilitate political organization and action, and provide a more decentralized environment for social, cultural and political interaction in which anybody can participate: User choice, application-agnosticism, innovation without permission, and low costs of application innovation. Network neutrality rules should not constrain the evolution of the network more than is necessary to reach these goals, make it easy to determine which behavior is and is not allowed to provide much-needed certainty for industry participants, and keep the costs of regulation low.

Second, over the years, the network neutrality debate has evolved into a series of sub-debates that are difficult to follow. As a result, even interested parties often lack a complete, accurate picture of the broader debate. This paper provides a comprehensive overview of existing proposals for non-discrimination rules from a variety of sources, including academics, industry participants, and policymakers in the US and abroad, and helps policy makers choose among the available options. It evaluates these proposals and the Open Internet Order's non-discrimination rule against the framework developed in *Part I* and proposes a non-discrimination rule that policy makers should adopt around the world – a rule that the Open Internet Order adopted in part.

In over ten years of debate, network neutrality proponents have struggled to come up with a rule that clearly specifies in advance which forms of differential treatment should be allowed. As a result, they had to fall back on all-or-nothing approaches or standards-based approaches, both of which create considerable social costs. The rule I propose – ban application-specific discrimination, allow application-agnostic discrimination – solves this problem. It accurately distinguishes between socially beneficial and socially harmful conduct (avoiding the problems of the all-or-nothing approaches), but does so *ex ante* (avoiding the social costs of the standards-based approaches).

[15] The following discussion draws in part on van Schewick (2012c).

Substantively, the rule balances the public interest in network neutrality with the legitimate interests of network providers. It prevents network providers from interfering with user choice or distorting competition among applications or classes of applications, while giving them broad flexibility to differentiate and price their Internet service offerings and manage their network in application-agnostic ways. The rule allows network providers to offer some forms of user-controlled Quality of Service and provides certainty to market participants. Technically, it reinforces key architectural principles on which the Internet was based without locking in the original architecture of the Internet itself.

The paper also explains why the following proposals for non-discrimination rules do not adequately protect the values that network neutrality rules are designed to protect: rules that allow all discrimination; rules that ban discrimination that violates an antitrust framework or ban behavior that is anticompetitive; case-by-case approaches that leave the decision over which discriminatory conduct should be banned to future adjudications; rules that ban discriminatory conduct that is not disclosed; rules that allow discrimination among applications or classes of applications that are not alike as long as the network provider does not discriminate among like applications or classes of applications. All of these proposals are currently under active consideration in the US or abroad.

Third, this paper exposes the deep disconnect between those who, like most network neutrality proponents and the FCC in the Open Internet Order, base calls for network neutrality regulation on a broad theoretical framework that considers a wide range of economic and non-economic harms, and those who evaluate calls for network neutrality regulation based on an antitrust framework. As this paper shows, these two frameworks lead to very different conclusions regarding which forms of differential treatment are problematic. Since these underlying theoretical differences are usually not made explicit, participants in the debate often talk past each other. For example, economists scold the FCC for proposing or adopting rules that are overreaching[16] – which is correct if you view the debate through an antitrust framework. In the context of the broad theoretical framework that the FCC explicitly adopts in its order, however, the Open Internet Rules make perfect sense. At the same time, network neutrality proponents who are not aware of the implications of the different frameworks do not necessarily realize that antitrust-based approaches or approaches that use terms that have well-defined meanings in antitrust law may reach correct results in the context of an antitrust framework, but do not capture many instances of discrimination that network neutrality proponents are concerned about. Thus, adopting such rules would make it impossible to successfully bring complaints against discriminatory conduct that violates the values that network neutrality rules are designed to protect.

Fourth, the paper offers the first in-depth analysis of the relationship between network neutrality and Quality of Service. This relationship is currently uncertain and contentious. Often, it is not immediately apparent how a specific non-discrimination rule affects network providers' ability to offer Quality of Service.

[16] See, e.g., Hazlett & Wright (2011), p. 6, parts V-VI.

The network neutrality debate is often framed as a debate for or against Quality of Service.[17,18] As the paper shows, the reality is much more subtle. Many network neutrality proposals allow some, but not all forms of Quality of Service, with different proposals drawing the line between acceptable and unacceptable forms of Quality of Service in different ways.[19] Underlying these differences are disagreements over the social desirability of different forms of Quality of Service. The paper advances the debate both descriptively, by examining which forms of Quality of Service would be allowed by which rule, and normatively, by exploring which, if any, forms of Quality of Service a network neutrality regime should allow. While many forms of Quality of Service allow Internet service providers to distort competition among applications and interfere with user choice, some forms of user-controlled Quality of Service do not. If a network neutrality regime includes certain restrictions on charging and provisions that protect the quality of the baseline service from dropping below unacceptable levels, these forms of Quality of Service provide the social benefits of different types of network service without the social costs, and should be allowed. The non-discrimination rule proposed by this paper and the Open Internet Order's non-discrimination rule allow network providers to offer these (and only these) forms of Quality of Service.[20]

Finally, the paper provides the first detailed analysis of the Open Internet Order's non-discrimination rule for fixed broadband Internet access and of its implications for network providers' ability to manage their networks and offer Quality of Service. Since the recent D.C. Circuit Court opinion has vacated the FCC's Open Internet Rules, policy makers in the US need to decide whether they should preserve the Open Internet Order's non-discrimination rule or adopt a different type of non-discrimination rule. Thus, an accurate understanding of the rule is central to the current debate over the future of network neutrality in the US. The rule bans discrimination that is unreasonable, subject to reasonable network management. Whether specific discriminatory conduct is unreasonable will be decided in future case-by-case adjudications. Thus, it is not immediately apparent which types of differential treatment the rule forbids. Drawing on a close reading of the text of the order, the paper sets out the Open Internet Order's non-discrimination standard as clarified by the text of the order and shows how it may apply to specific discriminatory conduct, in particular to the provision of Quality of Service.

[17] This perception is particularly common in Europe. *See, e.g.*, AT&T, et al. (2008), p. 2 ("'Net neutrality' advocates typically suggest . . . an 'all bits must be treated the same' approach."); Cave & Crocioni (2007), p. 671 (calling "no prioritisation" "one of the two main elements . . . of network neutrality"); Cave, et al. (2009), p. 1 ("The key issue is whether internet service providers should be prevented from introducing differentiated quality of service levels on the Internet."); Schuett (2010), p. 1 ("Roughly speaking, network neutrality refers to the principle that all data packets on an information network are treated equally.").

[18] Participants in the debate also often assume that the broad version of the end-to-end arguments would make it impossible to offer Quality of Service. This assumption is not correct, either. van Schewick (2010d), pp. 106-07.

[19] Quality of Service can be offered in different ways, which create different social benefits and social costs. For example, under some Quality of Service architectures, network providers decide which applications get which form of Quality of Service. Under others, users make that choice. Network providers may make Quality of Service available exclusively to individual applications or to classes of applications with similar needs.

[20] The paper's findings with respect to network neutrality and Quality of Service are summarized in more detail the Conclusion.

According to the text of the order, the FCC will evaluate discriminatory conduct under the non-discrimination rule and the reasonable network management exception based on how well they preserve two of the factors that were at the core of the Internet's succes: user choice and application-agnosticism.[21] This standard allows certain forms of user-controlled Quality of Service. The order explicitly rejects attempts to base non-discrimination rules on an antitrust framework. Banning only discrimination that violates the antitrust laws or is "anticompetitive," the order explains, would be too narrow and would not capture all instances of discrimination that the Open Internet Rules are concerned about. While the FCC did not adopt the non-discrimination rule proposed by this paper, the proposal heavily influenced the Open Internet Oder's non-discrimination rule. In particular, whether discriminatory behavior complies with the proposed rule (i.e. whether it is application-agnostic) is one of the factors the FCC will use to determine whether the conduct violates the FCC's non-discrimination rule and the reasonable network management exception. Thus, this paper's discussion of application-specific and application-agnostic discrimination can illuminate the rationale underlying the FCC's rule as well as help apply these provisions to specific instances of discriminatory conduct in the future.

Many network neutrality proponents were disappointed by the Open Internet Rules. While the rules are not perfect, the paper shows that they provide the FCC with a powerful set of tools to protect users and innovators against discrimination by providers of fixed broadband Internet access service.

This paper is part of the broader debate over network neutrality that has been raging for more than ten years. While the debate originally focused on the need for rules against blocking and discrimination, it has since evolved into a number of sub-debates. Each sub-debate focuses on a specific way in which a network provider could exploit its ability to control or interfere with the applications on its network and discusses whether rules are needed to address the problems this particular practice may cause.

As a result, the question "Should we adopt network neutrality rules?" can no longer be answered with a simple "yes" or "no." Instead, legislators and regulators considering whether to enact network neutrality rules need to answer a series of questions as they decide which, if any, network neutrality rules they should adopt (see *Box 1: Thinking About Network Neutrality Rules*). The first question is: "Do we need a rule against blocking, i.e. a rule that forbids network providers from blocking access to applications, content and services on their networks?" Such a rule is part of all network neutrality proposals; this is the one rule on which all network neutrality proponents agree. This paper assumes that the case for a rule against blocking has been made.[22]

This paper focuses on the second question: "Should the rules also ban differential treatment that falls short of blocking ("discrimination"), and, if so, which forms of differential treatment should be banned?" For example, if a network provider slows down Internet video

[21] In addition, preserving the freedom to innovate without permission is an explicit purpose of the Open Internet Rules. Thus, this factor can be used to interpret any provision of the rules, including the non-discrimination rule.

[22] See, e.g.,van Schewick (2010d), ch. 6-9.

applications like Netflix, Hulu or YouTube that compete with the network provider's own Internet video application, or provides low-delay service only to its own Internet video application, should these practices be prohibited?

The answer depends in part on the framework we use to evaluate network neutrality rules – whether we use an antitrust framework or the broader theoretical framework used by most network neutrality proponents and the FCC in its Open Internet Order. This paper assumes that the case for the broader theoretical framework has been made.

Box 1
Thinking About Network Neutrality Rules

Legislators and regulators considering whether to enact network neutrality rules need to answer the following questions as they decide which, if any, network neutrality rules they should adopt.

A. Goals of the Rules

General-Purpose Infrastructure vs. Antitrust

B. Rules

1. Ban blocking?
2. Ban discrimination? (This paper)
3. Impose limits on charging?
    a. for access to end users?
    b. for enhanced access to end users?

C. Exceptions

1. For what?
2. How to define?

D. Scope

1. Wireline vs. Wireless
2. Internet Access Services vs. Specialized Services

This paper does not address whether and, if so, whom network providers should be allowed to charge for any differential treatment that the chosen non-discrimination rule allows (see *Box 2: Charging for Quality of Service* below). Concerns about offering differential treatment and about charging for it are driven by different sets of policy considerations which should be considered and evaluated separately. I take up restrictions on charging elsewhere.[23]

[23] van Schewick (2010d), pp. 278-80, 290-93; van Schewick (2010h); van Schewick (2010a), pp. 10-12; van Schewick (2014b), Section 3; van Schewick (2014a), Section "Tough Lessons from Mobile and Music."

There, I argue that network providers should only be allowed to charge their own Internet service customers for any differential treatment allowed by the non-discrimination rule.[24]

Box 2
Charging for Quality of Service

> If policy makers adopt a non-discrimination rule that allows network providers to offer some form of Quality of Service or other forms of enhanced treatment, they need to decide whether, and if so, whom network providers should be allowed to charge for it. Again, policy makers have a number of options, each supported by at least some proponents of network neutrality: (1) the network provider is not allowed to charge anyone for the use of Quality of Service (though it can increase the general price for Internet service);[25] (2) it can charge only its Internet service customers;[26] (3) it can charge its Internet service customers and/or application and content providers, but is required to offer the service to application and content providers on a non-discriminatory basis;[27] (4) it can charge its Internet service customers and/or application and content providers.[28]

In sum, the paper assumes that the decision to base network neutrality regulation on a theoretical framework that considers a broader range of harms than an antitrust framework and the decision to adopt a rule against blocking have been made.

The paper proceeds in three parts. *Part I* sets out criteria that policy makers and others can use to choose among alternative proposals for network neutrality rules, interpret existing rules and evaluate specific forms of discriminatory conduct.

Part II evaluates eight existing proposals for non-discrimination rules against these criteria and proposes a non-discrimination rule that policy makers should adopt. In the process, it explains how the different non-discrimination rules affect network providers' ability to offer Quality of Service and which, if any, forms of Quality of Service a non-discrimination rule should allow.

---

[24] In addition, I argue that network neutrality proposals should ban any access charges to application and content providers (i.e., they should prohibit network providers from charging application or content providers who are not their Internet service customers for the right to access the network provider's Internet service customers), not just access charges in return for better transport. See van Schewick (2010h); van Schewick (2010a), pp. 7-10.

[25] E.g., Internet Freedom and Nondiscrimination Act of 2006 (2006), §3(2); Network Neutrality Act of 2006 (2006), §4(a)(7); Internet Freedom Preservation Act (2007), §12(a)(5).

[26] E.g., Lessig (2006), pp. 8–10. For criticism of this proposal, see Frischmann (2005), pp. 1009–1012, and Crawford (2007), pp. 403–404. Most proposals in this category would ban any access charges to application and content providers (i.e., they would prohibit network providers from charging application or content providers who are not their Internet service customers for the right to access the network provider's Internet service customers), not just access charges in return for better transport. See, e.g., Internet Non-Discrimination Act of 2006 (2006), §4(a)(3); Lee & Wu (2009); van Schewick (2010h); van Schewick (2010a), pp. 7-10; van Schewick (2014b), Section 3; van Schewick (2014a), Section "Tough Lessons from Mobile and Music." For criticism of these proposals, see Hahn & Wallsten (2006) and Hemphill (2008).

[27] E.g., Lessig (2008b), p. 9. For criticism of this proposal, see Sidak (2006).

[28] See, e.g., Verizon & Verizon Wireless (2007), pp. 42-47; Verizon & Verizon Wireless (2010), pp. 68-74.

Part III sets out the Open Internet Order's non-discrimination rule for fixed broadband Internet access. It evaluates the rule against the criteria used throughout this paper and discusses how the rule affects network providers' ability to offer Quality of Service.

Opponents of network neutrality regulation have created the impression that network neutrality rules force policy makers to choose between protecting users and application innovators against interference from network providers on the one hand and innovation in the network and the needs of network providers on the other hand. This paper refutes that myth. It shows how policy makers can protect users and innovators while also giving network providers the tools they need to manage their networks and allowing the network to evolve.

A FRAMEWORK FOR EVALUATING NETWORK NEUTRALITY RULES

When evaluating alternative proposals for non-discrimination rules or other network neutrality rules, legislators or regulators should consider a number of factors. Non-discrimination rules are part of a set of network neutrality rules that share common goals. Thus, an important criterion in evaluating proposals is how well they support these goals. The answer depends, of course, on what these goals are.

Some participants in the network neutrality debate view the debate through an antitrust lens. They interpret concerns about blocking, discrimination or other practices as concerns about anticompetitive vertical leveraging or vertical foreclosure, and apply an antitrust framework to evaluate and address these concerns.[29] Among network neutrality proponents, this is a minority position. Most network neutrality proponents base their calls for regulation on a theoretical framework that considers a wider range of economic and non-economic harms.[30] The FCC's Open Internet Rules are based on this broader framework as well.[31] Due to these differences, proponents of an antitrust framework and proponents of a broader framework will reach differing conclusions when evaluating proposals.[32] This paper assumes that the case for the broader theoretical framework has been made.

According to this broader theoretical framework, network neutrality regulation serves a number of goals. Most generally, network neutrality rules are intended to preserve the Internet's ability to serve as an open, general-purpose infrastructure that provides value to society over

[29] See references cited throughout Section "Case-by-Case Approaches: Ban Discrimination that Violates an Antitrust Framework" below.

[30] See, e.g., Lessig (2001b); Frischmann (2005), pp. 1012-1022; Wu (2006b); Frischmann & van Schewick (2007), pp. 423-428; Crawford (2009), pp. 916-917, 919; Crawford (2007), pp. 380-392. Center for Democracy & Technology (2010), pp. 8-9, 30; Free Press (2010), pp. 76-77, 134-136; Open Internet Coalition (2010), p. 76; Center for Media Justice, et al. (2010), pp. 23, 24-28.

[31] Federal Communications Commission (2010c), pp. 4-11, paras 11-19, pp. 45-46, para 78 and 47 C.F.R. §8.1.

[32] See, e.g., Frischmann & van Schewick (2007), pp. 426-428; Kang (2007), pp. 6-14. See Section "Case-by-Case Approaches: Ban Discrimination that Violates an Antitrust Framework" below.

time in various economic and non-economic ways.[33] More specifically, network neutrality rules aim to, first, foster innovation in applications.[34] Fostering application innovation is not only critical for economic growth.[35] A greater number of applications also increases the Internet's potential to create value in the social, cultural, and political domains.[36] Second, network neutrality rules are designed to protect users' ability to choose how they want to use the network, without interference from network providers. This ability to choose is fundamental if the Internet is to create maximum value, for users and for society.[37] Third, network neutrality rules aim to preserve the Internet's ability to improve democratic discourse, facilitate political organization and action and to provide a decentralized environment for social, cultural and political interaction in which anybody can participate.[38]

Network neutrality rules also have social costs:[39] They limit the evolution of the network's core and limit network providers' ability to realize all potential efficiency gains or optimize the network in favor of the applications of the day. They may reduce network providers' profits, and, like any regulatory regime, create costs of regulation that burden providers, users and society as a whole.

Thus, the decision to adopt network neutrality rules is based on a trade-off.[40] In a way, the first three costs are the price for a system that can evolve and support new applications in the future.[41] And while lower profits may to some degree reduce network providers' incentives to deploy more and better broadband networks, letting network providers block, discriminate or charge access fees removes the very features that were at the core of the Internet's success. Given that there are other ways to foster broadband deployment that are not similarly harmful, sacrificing the very aspects that drive the Internet's value seems too high a price to pay.[42] As

[33] See, e.g., van Schewick (2004); Wu (2006c), p. 16; Frischmann & van Schewick (2007), pp. 423-428; van Schewick (2010d), pp. 387-388. See also Federal Communications Commission (2010c), pp. 4-11, paras 11-19, pp. 45-46, para 78.

[34] See, e.g., Lessig (2001a), pp. 246-249; Wu (2003), pp. 145-156; Wu (2004), pp. 71-75, 80-84; van Schewick (2004); van Schewick (2007).

[35] See, e.g., van Schewick (2004), pp. 346-349; Wu (2004), pp. 80-84; van Schewick (2010d), pp. 356-359.

[36] van Schewick (2010d), pp. 359-361.

[37] van Schewick (2010d), pp. 361-364. See also footnote 54 below and accompanying text and *Box 4: Application-Agnostic vs. Application-Blind* below.

[38] See, e.g., van Schewick (2010d), pp. 359-362, 364-365; Herman (2006), pp. 112-119; Balkin (2008), pp. 101-107; Balkin (2009); Ammori (2009). On the Internet's social, cultural and political potential, see Benkler (2006).

[39] See, e.g., van Schewick (2010d), pp. 365-368. These costs feature, of course, prominently in the arguments made by network neutrality opponents as well. See, e.g., Yoo (2005); Farber, et al. (2007); Becker, Carlton & Sider (2010), p. 510-519; Faulhaber & Farber (2010b), pp.313-324; Faulhaber (2011), pp. 22-23.

[40] van Schewick (2010d), pp. 368-371; Frischmann & van Schewick (2007), pp. 423-425. For a detailed analysis of the trade-off from the perspective of a network neutrality opponent, see Yoo (2005), pp. 60-68, 70-76.

[41] van Schewick (2010d), pp.368-370.

[42] van Schewick (2010d), pp. 370-371.

Tim Wu put it, it is like selling the painting to get a better frame.[43] Based on this reasoning, network neutrality proponents resolve the trade-off in favor of the social benefits.[44,45]

Any network neutrality rule will impose these types of social costs, but different proposals for a non-discrimination rule will support the goals of network neutrality regulation to varying degrees and will have different social costs. Most generally, policy makers should choose the rule that realizes the goals of network neutrality regulation and imposes the least social costs.[46]

In line with these considerations, a non-discrimination rule (or any other network neutrality rule) should meet the following criteria:[47]

First, as I have explained elsewhere, there are a number of factors that have allowed the Internet to foster application innovation, improve democratic discourse, facilitate political organization and action, and provide a more decentralized environment for social, cultural and political interaction in which anybody can participate. They need to be preserved to allow the Internet to continue to do so in the future. These factors should serve as guiding principles not only when choosing among alternative options for network neutrality rules, but also when evaluating discriminatory conduct under existing network neutrality rules.[48,49] They are:[50]

- *Innovation without permission:* Innovators independently choose which applications they want to pursue; they do not need support or "permission" from network

[43] Wu (2006d).

[44] See, e.g., Wu (2006b), pp. 3-4; Frischmann & van Schewick (2007), pp. 423-425; van Schewick (2010d), pp. 368-371.

[45] Opponents of network neutrality regulation come to a different conclusion. In particular, they stress the potential impact of network neutrality rules on incentives to invest in network infrastructure (e.g., Yoo (2005), pp. 48-53; Becker, Carlton & Sider (2010), p. 506; Cisco Systems (2010), pp. 5-7).

[46] This is not a strict optimization problem. Not only may the different types of social costs be difficult to quantify exactly, they may also be incommensurable.

[47] See, e.g., van Schewick (2010e), p. 1 (asking the FCC to use these criteria in evaluating alternative proposals for non-discrimination rules in the context of the Open Internet proceeding); van Schewick (2010f), pp. 4-5 (same).

[48] See, e.g., van Schewick (2010i) (asking the FCC to choose network neutrality rules that preserve application-blindness, user choice, innovation without permission and low costs of application innovation); van Schewick (2010k), pp. 1-2 (arguing the FCC should use "these factors as guiding principles when choosing among alternative options for network neutrality rules, as well as when interpreting any network neutrality rules that should be adopted in the future.").

[49] This paragraph is adopted from van Schewick (2010k), pp. 1-2.

[50] The factors that have fostered application innovation in the past factors are described in detail in van Schewick (2010d). For a short overview, see van Schewick (2010i). For a brief discussion of the factors that are at the core of the Internet's political, social and cultural potential, see van Schewick (2010d), pp. 359-365; Benkler (2000), pp. 565-568; Balkin (2009). The original Internet created an environment characterized by these factors as a consequence of its architectural design. In particular, they are the result of the application of the layering principle and the broad version of the end-to-end arguments. On the layering principle, the broad version of the end-to-end arguments and their relationship to the original architecture of the Internet, see footnote 2 above and van Schewick (2010d), pp. 61-75, 96-103; van Schewick (2004). On early arguments that the architecture of the Internet, due to the end-to-end arguments, created a beneficial environment for innovation that regulation should preserve, see Lemley & Lessig (1999) (in the context of the debate over open access to cable networks) and, in the context of network neutrality, Lessig (2001a); Lessig (2002); Wu (2003); Wu & Lessig (2003); van Schewick (2004); Wu (2004); Cerf (2006); Lessig (2006); Lessig (2008b).

providers in order to realize their ideas for an application. Adding additional decision-makers who need to endorse the idea or take action before an idea can be realized reduces the chances that innovative ideas can be realized.[51]

- *User choice:* Users independently choose which applications they want to use, without interference from network providers.[52] Letting users, not network providers choose which applications will be successful is an important part of the mechanism that produces innovation under uncertainty.[53] At the same time, letting users choose how they want to use the network enables them to use the Internet in a way that creates more value for them (and for society) than if network providers made this choice.[54] (See *Box 3: The Importance of User Choice* below.)
- *Application-Agnosticism:* The network is application-agnostic. While an application-agnostic network may have information about the applications on the network, it does not make distinctions among data packets based on that information.[55] This ensures that network providers cannot interfere with innovators' and users' choices, that they cannot distort competition among applications (or classes of applications) or reduce application developers' profits through access fees[56] (we may call this *"innovation without fear"*). (On the relationship between application-agnosticism and applications-blindness, see *Box 4: Application-Agnostic v. Application-Blind* below.)
- *Low costs of application innovation:* The low costs of application innovation not only make many more applications worth pursuing, but also allow a large and diverse

[51] On innovation without permission in the original Internet, see van Schewick (2010d), pp. 204, 211, 293. On the impact of innovation without permission on innovation, see van Schewick (2010d), pp. 345-348. See also Cerf (2006), pp. 1,4; Balkin (2009) (focusing on the social, cultural and political implications).

[52] van Schewick (2010d), pp. 144, 152-155, 293-295, 362-364; Cerf (2006), pp. 1-3, 7.

[53] See van Schewick (2010d), pp. 349-351; van Schewick (2010i), p. 6; s*ee also* footnote 54 below.

[54] See van Schewick (2010d). pp. 362-363. See also Cerf (2006), pp. 1-3, 7. On the importance of user choice for the Internet's social, cultural and political potential, see, e.g., Balkin (2009); van Schewick (2010d), pp. 359-365.

[55] The original Internet was application-blind and application-agnostic. This was a consequence of its architecture, in particular of the broad version of the end-to-end arguments and of the layering principle. See *Box 4: Application-Agnostic v. Application-Blind* below, footnote 2 above and van Schewick (2010d), pp. 72-75, 217-218; van Schewick (2004). See also, e.g., Lemley & Lessig (1999) , para 17, Cerf (2006), pp. 1-4, 7; Reed (2010). For a short summary of the importance of application-blindness, see van Schewick (2010i), pp. 3-4. For a detailed analysis, see van Schewick (2010d), pp. 215-281, 286-295, 349-353, 355-365. While the analysis in these sources focuses on the impact of application-blindness, the analysis equally applies to application-agnosticism. An application-blind network is necessarily application-agnostic. In particular, both create the same environment for application innovation and network use (see *Box 4: Application-Agnostic v. Application-Blind* below*).* Thus, their economic, social, cultural and political impact is the same. See also Benkler (2000), pp. 565-568; Balkin (2009); van Schewick (2010d), pp. 359-365 (all focusing on the social, cultural and political implications).

[56] Access fees are fees that the network provider imposes on application and content providers who are not its Internet service customers. Access fees come in two variants: In the first variant, a network provider charges application or content providers for the right to access the network provider's Internet service customers. In the second variant, which is sometimes called "paid prioritization" or "third-party-paid prioritization," a network provider charges application or content providers for prioritized or otherwise enhanced access (e.g., access that does not count towards the users' monthly bandwidth cap) to these customers. A discussion of access fees is outside the scope of this paper. On access fees, see, e.g., van Schewick (2010h); see also *Box 2: Charging for Quality of Service* above and footnotes 23-24 above and accompanying text.

group of people to become innovators.[57] If there is uncertainty (e.g., about technology or user needs) or user needs are heterogeneous, a larger and more diverse group of innovators will create more and better application innovation than a smaller, less diverse group of innovators, and these applications will better meet the needs of Internet users.[58] In the current Internet, there is uncertainty and user needs are heterogeneous, so the conditions under which innovator diversity increases the amount and quality of innovation are met.[59]

Second, a non-discrimination rule should not constrain the evolution of the network more than is necessary to reach the goals of network neutrality regulation and not impose other unnecessary social costs.

Finally, the rule should make it easy to determine which behavior is and is not allowed to provide certainty for industry participants. For application providers, uncertainty over the level of protection provided by the rule reduces their incentives to innovate and their ability to get funding.[60] For network providers, uncertainty over the legality of network management practices or of different forms of Quality of Service may make it more difficult to manage their network and may limit the evolution of the network infrastructure. Uncertainty over the regulatory regime may also reduce network providers' incentives to invest more generally.[61] Thus, certainty increases the social benefits and reduces the social costs of a non-discrimination rule.

In sum, policy makers should look for a rule that fosters application innovation, protects user choice and preserves the Internet's economic, social, cultural and political potential while avoiding unnecessary social costs (for a summary of the framework, see *Box 5: A Framework for Evaluating Network Neutrality Rules* below). In particular, the rule should preserve user choice, innovation without permission, application-agnosticism and low costs of application innovation. As this paper will show, this framework can not only help guide the choice among alternative proposals for non-discrimination rules or other network neutrality rules. It can also be used to interpret existing non-discrimination rules or evaluate specific forms of discriminatory conduct.

[57] For a short version of the argument, see van Schewick (2010i), pp. 2-3, 5-6 and van Schewick (2010h), pp. 4-5. On low cost of application innovation in the original Internet, see van Schewick (2010d), pp. 138-148, 204-205, 289-290. On the impact of low cost innovation on who can innovate, see van Schewick (2010d), pp. 204-213, 292-293. See also Benkler (2000), pp. 565-568; Balkin (2009) (both focusing on the social, cultural and political implications).

[58] For a short version of the argument, see van Schewick (2010i), pp. 5-6 and van Schewick (2010h), pp. 4-5. For the detailed version, van Schewick (2010d), pp. 298-349.

[59] See van Schewick (2010d), p. 356.

[60] See footnotes 229 to 231 below and accompanying text

[61] See footnotes 226 to 228 below and accompanying text.

Box 3

The Importance of User Choice

Why the emphasis on user choice? As I explained in my testimony at the Federal Communications Commission Second En Banc Hearing on Broadband Management Practices in April 2008, "[f]irst, user choice is fundamental if the Internet is to create the maximum value to society. The Internet is a general purpose technology. It does not create value through its existence alone. It creates value by enabling users to do the things they want or need to do. Users know best what this is. As a result, users, not network providers should be able to decide how they would like to use the network, and what is important to them. Of course, in order for users to behave efficiently, they also need to bear (at least some of) the costs of their actions, something which the current system does not sufficiently provide.

User choice is also a fundamental component of the mechanism that enables application-level innovation to function effectively. In the current Internet, it is impossible to predict what future applications will be successful. Enabling widespread experimentation at the application-level and enabling users to choose the applications they prefer is at the heart of the mechanism that enables innovation under uncertainty to be successful.

By singling out specific applications, network providers pick winners and losers on the Internet. Whom they pick may be driven by a number of motivations that are not necessarily identical with what users would prefer, leading to applications that users would not have chosen and forcing users to engage in an Internet usage that does not create the value it could. Consumers, not network providers, should continue to choose winners and losers on the Internet."[62]

[62] Taken, with slight adaptations, from van Schewick (2008a), pp. 7-8.

Box 4

Application-Agnostic vs. Application-Blind

The Internet's original architecture was based on the layering principle and on the broad version of the end-to-end arguments. As a consequence of that design, the Internet was application-blind and application-agnostic. An application-blind network is unable to distinguish among the applications on the network, and, as a result, it is unable to make distinctions among data packets based on this information.[63] Unlike an application-blind network, an application-agnostic network may have information about the applications on the network, but, like an application-blind network, it does not make distinctions among data packets based on this information.

Thus, an application-blind network is necessarily application-agnostic: it does not make distinctions among data packets based on information about the applications on the network, because it does not have this information. By contrast, an application-agnostic network is not necessarily application-blind, because it may have information about the applications on the network.

If the Internet's original architecture was both application-blind and application-agnostic, why does the paper require the network to be application-agnostic and not application-blind?

For network providers, information about the applications on their network may be useful for capacity planning or security. Data on patterns of network use may enable network providers to predict or at least observe changes in the behavior of users, which may facilitate capacity planning.[64] Similarly, a clear overview of the applications using a network at a specific point in time may make it easier to detect security attacks.[65] At the same time, the potential harm to application innovation and user choice does not arise from information about the use of the network as such, but from network providers' ability to use that information to distort competition among applications or classes of applications or to interfere with user choice, e.g., by using this information to block, discriminate against or charge for the use of specific applications or classes of applications. Thus, from a policy perspective, the focus on application-agnosticism balances the public interest in protecting users and application providers from interference from network providers on the one hand and the needs of network providers on the other hand. From an architectural perspective, focusing on application-agnosticism instead of application-blindness constitutes a trade-off between architectural purity and the interests of network providers in being able to collect information about the use of the network.[66]

(continued)

[63] See footnote 2 above and van Schewick (2010d), pp. 72-75, 217-218; van Schewick (2004). See also, e.g., Cerf (2006), pp. 1-4, 7; Reed (2010).

[64] Blumenthal & Clark (2001), pp. 77-78, 86; St. Johns & Huston (2003), Section 3; Clark, et al. (2005), p. 472. This sentence and the next draw on van Schewick (2010d), p. 73.

[65] Blumenthal & Clark (2001), pp. 80; St. Johns & Huston (2003), Section 4.

[66] Using devices in the network to access higher-layer protocol data units (or, less technically, the messages passed to the Internet layer by a higher layer for delivery to its higher layer protocol peer) to gather information violates the layering principle. However, as long as the device does not modify or act on that higher-layer data, the independence of lower layers from higher layers, one of the key features of layering, is still preserved. Thus, the negative effect of this violation will be marginal. By contrast, the architectural effect of devices that modify or act upon information contained in higher-layer protocol data units is usually quite severe. For a longer explanation, see van Schewick (2010d), pp. 74-75.

This analysis is not meant to imply that network providers should have the unlimited right to collect information about applications or user behavior. The erosion of application-blindness in today's Internet not only threatens the Internet's ability to reach its economic, social, cultural and political potential; it also threatens users' privacy.[67] Network neutrality rules address the first threat only; they are not concerned with user privacy. While application-agnosticism adequately protects the values that network neutrality is designed to protect, privacy values may require stronger limits on the visibility of information in the network. Thus, even if network neutrality rules do not restrict network providers' ability to collect information about applications or user behavior, privacy law may need to limit that ability to address the privacy threats arising from the erosion of application-blindness in today's Internet.[68,69]

In its Open Internet order, the FCC adopted a similar framework. The rules are explicitly based on the broader theoretical framework supported by most network neutrality proponents and used by this paper.[70] Preserving user choice and innovation without permission is an explicit purpose of the rules. [71] Thus, these factors can be used to guide the interpretation of individual provisions of the Open Internet Rules. Moreover, as will be explained in more detail below,[72] the text of the order explicitly specifies that the FCC will evaluate discriminatory conduct based on how well it preserves user choice and the application-agnosticism of the network in order to determine whether the behavior is "reasonable" and therefore allowed under the Open Internet Rules' non-discrimination rule. (The same factors will be used to evaluate discriminatory or exclusionary conduct under the rules' exception for reasonable network management.) Thus, the FCC explicitly specifies that it will interpret key provisions of its rules – the non-discrimination rule for fixed broadband Internet access and the reasonable-network-management exception – based on how well they preserve three of the factors used to evaluate alternative options for non-discrimination rules and specific discriminatory conduct throughout this paper: user choice, application-agnosticism and innovation without permission.[73] Thus, the results of this paper indicate how the non-discrimination rule and reasonable-network-

[67] Reed (2008).

[68] By contrast, if network neutrality rules prohibited network providers from collecting information about the applications on their networks, these rules would have the side-effect of safeguarding users' privacy. Similarly, strong privacy laws that prohibit network providers from gathering that information would also re-create the application-blindness of the network, making it impossible for network providers to block, discriminate or charge differently based on that information. On the interactions between network neutrality laws and privacy laws, see also Ohm (2009), pp. 1489-1496.

[69] How privacy law should react to the erosion of application-blindness in today's Internet is beyond the scope of this paper. On this question, see, e.g., Ohm (2009); Cooper (2011).

[70] Federal Communications Commission (2010c), pp. 4-11, paras 11-19, pp. 45-46, para 78 and 47 C.F.R. §8.1.

[71] 47 C.F.R. § 8.1 (2011) ("The purpose of this Part is to preserve the Internet as an open platform enabling *consumer choice*, freedom of expression, *end-user control*, competition, and the freedom to *innovate without permission*.") (emphasis added).

[72] See the discussion of the FCC's non-discrimination rule in Part III of this paper.

[73] Since preserving innovation without permission is an explicit purpose of the Open Internet Rules, this factor can be used to interpret any provision of the Open Internet rule, including the non-discrimination rule.

management exception could apply to specific discriminatory conduct, which, in turn, may help adjudicators apply these rules in specific cases or help market participants understand the implications of these rules in more detail.

Box 5
A Framework for Evaluating Network Neutrality Rules

A network neutrality rule should meet the following criteria:

- It should preserve the factors that have allowed the Internet to serve as a platform for application innovation, free speech and decentralized economic, social, cultural and political interaction in the past:
  - *User choice:* Users independently choose which applications they want to use, without interference from network providers.
  - *Innovation without permission:* Innovators independently choose which applications they want to pursue; they do not need support or "permission" from network providers in order to realize their ideas for an application.
  - *Application-agnosticism:* The network is application-agnostic. While an application-agnostic network may have information about the applications on the network, it does not make distinctions among data packets based on this information.
  - *Low costs of application innovation:* The costs of application innovation are low.
- It should not constrain the evolution of the network more than is necessary to reach the goals of network neutrality regulation.
- It should make it easy to determine which behavior is and is not allowed to provide much-needed certainty for industry participants.
- It should keep the costs of regulation low.

PROPOSALS FOR NON-DISCRIMINATION RULES

When determining whether to adopt network neutrality rules, legislators and regulators need to decide whether the network neutrality rules should only ban blocking, or also discrimination, and if they decide to ban discrimination, how discrimination should be defined. The answers to these questions may affect how the core of the network can evolve. In particular, they determine whether a network provider can offer Quality of Service.[74]

[74] On Quality of Service, see footnotes 10-14 above.

SCOPE OF NON-DISCRIMINATION RULES

Non-discrimination rules apply to any form of differential treatment that falls short of blocking. The most obvious examples involve differential handling of data packets associated with different applications or uses. For example, a network provider may provide a low-delay service to its own streaming video application, but not to competing streaming video applications. Streaming video applications are sensitive to delay, so this increases the relative performance of the network provider's own application during times of congestion.[75]

The scope of non-discrimination rules is, however, not restricted to differential handling of packets in the network. Network neutrality rules aim to prevent network providers from distorting the playing field among applications or classes of applications, and from interfering with users' choices regarding the use of the network. In line with this goal, non-discrimination rules apply to any form of differential treatment that may make some applications or classes of applications or uses relatively more attractive to users than others. Consider an Internet service provider whose Internet service is subject to a monthly bandwidth cap. Traffic associated with its own streaming video application does not count towards the bandwidth cap, while the traffic of unaffiliated applications counts towards the bandwidth cap. (For an example, see *Box 6: Differential Counting of Traffic Towards the Monthly Bandwidth Cap: Comcast's Xfinity TV App for the Xbox* below). Other things being equal, this makes the Internet service provider's streaming video application relatively more attractive. For a video that produces 2 GB of traffic, using an unaffiliated application brings the user 2 GB closer to his bandwidth cap. By contrast, a user who chooses the Internet service provider's application keeps the option of using the 2 GB of his monthly bandwidth allowance that would have been used by this video for something else instead. Users of the unaffiliated application who have exhausted the monthly bandwidth allowance will have to bear the consequences of using another 2 GB (e.g., paying overage charges, having their traffic throttled or being cut off from Internet access), while users of the affiliated application do not face any consequences. Thus, although the data packets associated with different streaming video applications receive the same technical treatment in the network,[76] the practice of counting only some streaming video applications towards the monthly bandwidth cap would still be subject to the non-discrimination rules discussed in this paper.

[75] The size of the advantage is related to the size of delay and how well the different applications can cope with increases in delay.

[76] It is unclear whether Comcast also prioritizes traffic associated with the Xfinity TV App for the Xbox over other traffic travelling to and from the Xbox. See *Box 6: Differential Counting of Traffic Towards the Monthly Bandwidth Cap: Comcast's Xfinity TV App for the Xbox* below.

Box 6

Differential Counting of Traffic Towards the Monthly Bandwidth Cap: Comcast's Xfinity TV App for the Xbox

In Spring 2012, Comcast, the largest provider of Internet service in the US, introduced a new application for the Xbox, the Xfinity TV App. The Xfinity TV App allows Comcast's Internet service subscribers to view selected video content from Comcast's on-demand service on the Xbox, if they also subscribe to Comcast's traditional video offering and Microsoft Xbox Live Gold subscription service. Comcast's Internet service has a 250 GB monthly bandwidth cap.[77] Traffic associated with the Xfinity TV App to the Xbox does not count towards that cap, while traffic of other applications that also allow users to view on-demand video content on the Xbox (e.g., HBO Go or Netflix) does.[78] As explained in the text, this type of differential counting of traffic towards the monthly bandwidth cap is generally subject to the non-discrimination rules discussed in this paper, even if the data packets associated with the different streaming video applications receive the same technical treatment in the network.[79] Comcast claims, however, that the Xfinity TV App is not provided over the public Internet and is therefore not subject to the FCC's Open Internet Rules.[80] Whether this assessment is correct is outside the scope of this paper.

It is unclear whether Comcast also prioritizes traffic associated with the Xfinity TV App for the Xbox over other traffic travelling to and from the Xbox. Tests have shown, and Comcast admits, that Xfinity TV App traffic receives special marking. In tests, the marked Xfinity TV App traffic consistently outperformed unmarked Netflix traffic to the Xbox during times of congestion, but Comcast claims it is not prioritizing this traffic.[81]

Internet service providers may also favor specific applications or classes of applications over others by charging their subscribers different Internet-access fees for different applications or types of applications.[82] For example, an Internet service provider may charge a higher bandwidth-adjusted price for Internet access for Internet-telephony traffic than for the traffic of other applications, e.g., to extract more of the value that Internet telephony users realize from the use of that application.[83] Other things being equal, this increases the costs of using Internet

[77] Comcast Corporation (2012a).

[78] Comcast Corporation (2012b).

[79] For further discussion, see footnote 370 and *Box 18: Examples of Differential Impact: Comcast's Digital Voice service and Comcast's Xfinity TV App for the Xbox* below.

[80] Werner (2012).

[81] Dugan (2012); Higginbotham (2012); Werner (2012)).

[82] Currently, Internet service providers charge the same per-byte (or otherwise bandwidth-adjusted) price for all applications. For a detailed analysis of network providers' incentives to engage in application-specific pricing and of the impact on application developers and users, see van Schewick (2010d), pp. 273-375, 277-278. See also footnote 420 below.

[83] For a real-world example, see the presentation by two providers of deep packet inspection technology for mobile networks (Allot Communications & Openet (2010), p. 7), which describes a pricing scheme in which subscribers would pay 2c per MB for Facebook traffic, 3$ per month for Skype traffic, 50c per month for YouTube traffic und nothing for traffic associated with the mobile provider's content offerings.

telephony relative to other applications, making Internet telephony relatively less attractive.[84] Thus, Internet-access plans that charge different bandwidth-adjusted prices for different applications would be subject to the non-discrimination rules in this paper, even if the data packets associated with the different applications receive the same treatment in the network.

ALL- OR NOTHING-APPROACHES

A first set of approaches takes an all-or-nothing position towards differential treatment of packets.

Allow All Discrimination (or "No Rule Against Discrimination")

Network providers and other opponents of network neutrality regulation oppose any restrictions on network providers' ability to differentiate among data packets.[85] They would not adopt a non-discrimination rule and allow all discrimination instead. This approach would not restrict network providers' ability to offer Quality of Service in any way. The Federal Communications Commission's Open Internet Order adopted this approach for wireless networks in December 2010. The Open Internet Rules for wireless networks ban blocking of some, though not all applications, but do not impose any restrictions on discrimination.[86]

Proponents of this approach contend that a ban on all discrimination would make it impossible to offer services such as Quality of Service, or to manage networks during times of congestion.[87] Without Quality of Service, proponents of this option argue, certain types of applications – those that require special treatment from the network – will not be able to operate. Thus, banning Quality of Service may reduce innovation in applications that need or benefit from Quality of Service and deprive users and society of the benefits these applications would have created.[88] While some forms of differential treatments such as those involved in Quality of Service would be socially beneficial, the argument continues, trying to distinguish between beneficial and harmful discrimination (to the extent it exists) would be too difficult.[89] Since technology is evolving rapidly, regulators are likely to get it wrong.[90] Even if legislators or regulators succeed in identifying criteria that accurately distinguish between beneficial and

[84] On the potential harms resulting from differentiating among classes of applications, see footnotes 362 to 402 and accompanying text.

[85] In the US, see, e.g., AT&T Inc. (2007), p. iii; Clark, Lehr & Bauer (2010), pp. 21-22, 25; Cisco Systems (2010), pp. 5-8; Faulhaber & Farber (2010b), pp. 316-317, 336; Verizon & Verizon Wireless (2007), pp. 41-42, 47-56; National Cable & Telecommunications Association (2007), p. 36; in Europe, see, e.g., AT&T, et al. (2008), pp. 1-4; Cave, et al. (2009).

[86] Federal Communications Commission (2010c), pp. 52-58, paras 93-106.

[87] AT&T Inc. (2007), pp. 52-53.

[88] Litan & Singer (2007), pp. 569-570; Becker, Carlton & Sider (2010), pp. 511-512, 518-519; Faulhaber (2011), p. 13. Allowing Internet service providers to offer Quality of Service may also allow them to differentiate themselves from their competitors. Yoo (2005), pp. 29-33 (discussing the benefits of allowing network providers to differentiate their Internet service offerings in general); AT&T Inc. (2007), pp. 71-73; Verizon & Verizon Wireless (2007), pp. 57-58; TELUS Communications Company (2009), p. 20, para 58.

[89] AT&T Inc. (2007), pp. 83-85.

[90] Yoo (2005), p. 67; TELUS Communications Company (2009), p. 4, para 17; Becker, Carlton & Sider (2010), p. 509.

harmful discrimination when the regulation is enacted, these criteria may not be accurate in the future. For these reasons, regulators should give up on trying to separate socially beneficial from socially harmful discrimination and allow all discrimination instead.

This view fails to recognize that banning blocking, but allowing discrimination will make the rule against blocking meaningless by offering a legal alternative to blocking that is less costly and potentially more effective.[91] Discrimination is an attractive alternative to blocking. Discrimination reduces the perceived quality of the affected application relative to others. If a network provider secretly slows down packets or uses methods that are difficult to detect, its customers may attribute poor performance of the affected application or website to design flaws, and happily switch to the network provider's supposedly superior offering. While the result of blocking and discrimination is the same - the network provider's Internet service customers stop using the blocked or degraded application and switch to the application that is not blocked or degraded -, the costs of using discrimination instead of blocking are much lower. If the network provider blocks an application, users will notice and may switch to another Internet service provider.[92] By contrast, users who do not realize that their network provider interfered with their preferred application and think they chose the better application will have no incentive to switch.[93]

Based on these considerations, we would expect network providers to prefer discrimination over outright blocking. Network providers' actual behavior is in line with these predictions. In the examples that are often highlighted in the debate, network providers often use methods that make it more difficult or costly to reach particular applications or content instead of blocking access to them completely. For example, as the investigation of network providers' Internet traffic management practices by the Canadian Radio-television and Telecommunications Commission showed, most of the larger Canadian Internet service providers only throttled or otherwise interfered with peer-to-peer file-sharing applications, but did not block them completely.[94] Similarly, in 2009, British Telecom (BT) restricted the bandwidth available to the BBC iPlayer and other streaming video applications to 986 kilobytes/sec in BT's "Up to 8 Mbps Option 1" broadband service; it did not block these applications completely.[95] Available evidence suggests that network providers are well aware of the advantages of this strategy. In 2007, Comcast was found to be interfering with BitTorrent and other peer-to-peer file-sharing applications. To shut down BitTorrent connections, Comcast used "forged" data packets that seemed to come from the other party involved in the specific BitTorrent

[91] The following paragraph draws on van Schewick (2010d), pp. 260-261.

[92] A user who notices that his application is being blocked will not necessarily switch providers, even if she would have preferred to use the blocked application over alternative applications that are not blocked. See the discussion in Section "Ban Discrimination That Is Not Disclosed" below.

[93] Requiring network providers to disclose whether they interfere with applications and content will not fully solve this problem. See the discussion in Section "Ban Discrimination That Is Not Disclosed" below.

[94] Parsons (2009).

[95] Cellan-Jones (2009).

connection.[96] RCN, a competitive cable provider in the US, used the same technology from 2005 to 2009.[97] As white papers produced by Comcast's equipment vendor, Sandvine, showed, this method of interference was deliberately chosen to prevent customers from noticing the interference.[98] Network providers know that the use of file-sharing applications is an important driver of broadband adoption, and they do not want to lose customers who wish to use these applications.[99]

Finally, proponents of this approach implicitly assume that all forms of Quality of Service are equally beneficial. This assumption is not correct. Different forms of Quality of Service have different social benefits and social costs. Some are socially beneficial, some are socially harmful. As a result, a blanket permission of Quality of Service is not justified.

In sum, allowing all forms of discrimination does not adequately protect users and application developers against socially harmful discrimination and makes the rule against blocking meaningless. Thus, an effective network neutrality regime needs to ban blocking *and* socially harmful discrimination. As will be explained below, beneficial forms of discrimination can be accommodated through the definition of discrimination or through exceptions.

Ban All Discrimination

By contrast, some participants in the debate would ban all discrimination, requiring network providers to treat every packet the same.[100] The FCC's draft rules in the Open Internet Proceeding are an example of this type of non-discrimination rule.[101] A rule that required

[96] Comcast Corporation (2008a).

[97] RCN Corporation (2010).

[98] Sandvine (2004), p. 14.

[99] Mennecke (2005); Hellweg (2003); Sandvine (2004), pp. 5-6. See also National Cable & Telecommunications Association (2007), p. 31 ("Cable operators will not go down the path of blocking access to video or P2P services. Blocking such services would be recipe for […] massive dissatisfaction among consumers, which would lead to loss of customers to our competitors.")

[100] Internet Non-Discrimination Act of 2006 (2006); Crawford (2007), pp. 403-404. See also Frischmann (2010); Frischmann (2010); Frischmann (2012), pp. 348-355. While Frischmann's proposed non-discrimination rule would not require network providers to treat all packets the same (it would ban discrimination based on the identity of the user or use, whether or not there is congestion), Frischmann would ban all forms of Quality of Service or prioritization, even during times of congestion. See also Frischmann (2005), p. 1011-1012.

[101] Federal Communications Commission (2009b), paras 103-109, pp. 41-43 ("Subject to reasonable network management, a provider of broadband Internet access service must treat lawful content, applications, and services in a nondiscriminatory manner." ibid., p. 41, para 104, and noting that the proposed non-discrimination rule "bears more resemblance to unqualified prohibitions on discrimination added to Title II in the 1996 Telecommunications Act than it does to the general prohibition on "unjust or unreasonable discrimination" by common carriers in section 202(a) of the Act." ibid., p. 43, para 109). This rule would have been subject to an exception for reasonable network management. Federal Communications Commission (2009b), pp. 50-52, paras 135-141. This non-discrimination rule was supported by, e.g., Free Press (2010), pp. 74-75; Open Internet Coalition (2010), pp.15-17; Center for Media Justice, et al. (2010). These commenters would have coupled the non-discrimination rule with a relatively narrow reasonable network management exception. While the details of the proposed standards for defining "reasonable network management" differ, these commenters generally agreed that a particular practice would have to (a) serve a legitimate purpose related to the technical management of the network and (b) be narrowly tailored to address that purpose. See, e.g., Free Press (2010), pp. 78-79, 82-104; Open Internet Coalition (2010), pp. 41-50: Center for Media Justice, et al. (2010), p. 35-44.

network providers to treat every packet the same would make it impossible to offer Quality of Service, which, by definition, entails the network treating packets differently.[102]

Proponents of this option are concerned that network providers may use the provision of Quality of Service as a tool to distort competition among applications or classes of applications. For example, they are concerned that a network provider may offer Quality of Service exclusively to its own application, but not to other, competing applications, or may sell Quality of Service exclusively to one of several competing applications.[103] They also point out that network providers who offer Quality of Service and are allowed to charge for it have an incentive to reduce the quality of the baseline service below acceptable levels to motivate users to pay for better service.[104] Moreover, selling Quality of Service allows network providers to profit from bandwidth scarcity, which reduces their incentives to increase the capacity of their networks.[105] While these arguments all have merit, these problems can be solved without totally banning Quality of Service. As will be explained below, it is sufficient to constrain how Quality of Service can be offered and charged for.[106]

Supporters of banning Quality of Service also question the need for Quality of Service.[107] If there is no need for Quality of Service, banning it creates limited social costs.[108] So far, proponents of a ban point out, the lack of Quality of Service has not prevented real-time applications from becoming successful on the public Internet.[109] For example, although Internet telephony is sensitive to delay and high variations in delay ("jitter") and may benefit from a network service that provides low delay and low jitter, Internet telephony applications such as Skype or Vonage work in the current Internet.[110] Video telephony applications like Skype or

[102] Non-discrimination rules usually have an exception for reasonable network management. Thus, even under the strict non-discrimination rule described in the text, network providers may still be able to provide some or all forms of Quality of Service, if the form of Quality of Service under consideration meets the definition of reasonable network management. This, in turn, depends on the definition and interpretation of reasonable network management.

[103] Free Press (2010), pp. 21-23; Center for Media Justice, et al. (2010), p. 48, 51.

[104] See, e.g., Economides (2010), p. 8.

[105] Center for Democracy & Technology (2010), pp. 28-29; Economides (2010), pp. 6, 13; Free Press (2010), p. 22; Open Internet Coalition (2010), p. 46; Center for Media Justice, et al. (2010), p. 45.

[106] First, the non-discrimination rule I propose below allows only certain forms of Quality of Service. The constraints imposed by the rule make it impossible for network providers to use the provision of Quality of Service to distort competition among applications or classes of applications. (See the discussion on pp. 87-98 below and footnote 480 and accompanying text.) Second, my proposal requires the regulatory agency in charge of enforcing network neutrality rules to monitor the quality of the baseline service and set minimum quality standards, if the quality of the baseline service drops below acceptable levels. This prevents Internet service providers from degrading the quality of the baseline, best-effort service (e.g., by allocating less bandwidth to the best-effort service or by refraining from adding needed network capacity) to motivate users to pay for an enhanced type of service. (See footnote 479 and accompanying text and van Schewick (2010a), pp. 10-11.) Third, the proposed rule constrains how network providers can charge for Quality of Service. These constraints prevent network providers from charging in ways that would distort competition or harm application innovation. (See footnote 473 and accompanying text and van Schewick (2010h); van Schewick (2010i); van Schewick (2010a), pp. 10-12.)

[107] Center for Media Justice, et al. (2010), p. 49-50; Open Internet Coalition (2010), pp. 33-35.

[108] For a similar interpretation, see Clark, Lehr & Bauer (2010), p.10.

[109] Center for Media Justice, et al. (2010), p. 49-50; Open Internet Coalition (2010), pp. 33-35.

[110] Peterson & Davie (2012), p. 531.

Google Video Chat function over today's broadband connections.[111] The success of real-time applications on today's best-efforts Internet is due to two reasons: Many regions currently seem to have sufficient network capacity to prevent the lack of Quality of Service from becoming a problem.[112] If there is enough capacity so that congestion is generally low, the level of delay will be low enough most of the time to be tolerable for real-time applications.[113] In addition, network engineers and application designers have developed end-host-based techniques that allow real-time applications to compensate for the lack of Quality of Service in the network.[114] Pointing to this experience, proponents of a ban argue that capacity increases, combined with end-host based measures, are sufficient to meet the needs of applications that require low delay or low jitter.[115]

While available capacity affects the benefits of offering Quality of Service, the relationship between the two is more nuanced than is often assumed. Applications that would benefit from Quality of Service ("QoS-sensitive applications") are sensitive to the increase in delay, jitter or loss or to the variation in throughput that arises if queues build up in routers along the application's path, creating congestion (see *Box 7: The Relationship between Congestion, Delay, Jitter and Loss* below).[116] A network that offers Quality of Service can "help" these applications by providing classes of service that may offer throughput, delay, loss or jitter that are better suited to the needs of QoS-sensitive applications than the unpredictable and potentially highly variable throughput, delay, loss and jitter offered by the best-effort service.[117] Potential classes of service may offer throughput, loss, delay or jitter that is relatively better than the throughput, loss, delay or jitter provided by best-effort service during times of congestion[118] or may provide a performance that is more constant and predictable than best-effort service.[119,120] These services, however, can improve on the performance of best-effort service

[111] For example, Skype video requires a high-speed broad connection of at least 512kbps down / 128kbps up. For best quality, Skype recommends "a high-speed broadband connection of 4Mbps down / 512kbps up". Skype (2012).

[112] Davie (2003), p. 133.

[113] Kurose & Ross (2010), pp. 629-630. For a more detailed analysis of the relationship between capacity, congestion and Quality of Service, see the discussion in the following paragraphs.

[114] Kurose & Ross (2010), pp. 616-629.

[115] See, e.g., Open Internet Coalition (2010), pp. 43-46; Frischmann (2012), pp. 353-355

[116] The definition of congestion used throughout this section is derived from the definition of congestion in queuing theory. By contrast, network providers often define congestion differently. For more on the two definitions and their implications, see *Box 8: Definitions of Congestion and Benefits from Quality of Service* below.

[117] Different QoS-sensitive applications will often have different requirements in terms of throughput, delay, jitter or loss, so different QoS-sensitive applications may benefit from classes with differing characteristics. The combinations of throughput, delay, jitter and loss that could be provided by a class of service are limited. In any queuing system with finite buffers, there is a relationship between the distributions of loading factor, loss and delay. Combinations of two of these three variables determine the value of the third. In particular, "for fixed loss rate, reducing delay implies that throughput will fall. […] For fixed throughput, reducing delay implies an increase in loss rate. […] For fixed delay, reducing loss rate will reduce available throughput." See Davies, Holyer & Thompson (1999), pp. 2-3. The variability of delay in turn determines jitter.

[118] These types of service provide service that is as good as best-effort service if there is no congestion, and better than best-effort service if there is congestion. Huston (2012), p. 2.

[119] The performance of these services does not vary with congestion. As a result, their performance may be better than best-effort during times of congestion, but worse than best-effort if there is no congestion. This may occur, for

only if there is congestion.[121] If there is no congestion (i.e. if all queues are empty), congestion-related loss and queuing delay will be zero constantly and jitter will be low for all packets, and data flows will experience the maximum throughput and minimum end-to-end delay that is possible on their path.[122] No class of service can improve on that. Thus, Quality of Service is only useful if there is at least some congestion.

Box 7

The Relationship between Congestion, Delay, Jitter and Loss

Throughout this section, "congestion" denotes the building up of a queue for an outgoing link at a router, which may increase delay, jitter or packet loss.[123] (This definition differs from the definition of congestion that is often used by network providers. See *Box 8: Definitions of Congestion and Benefits from Quality of Service* below.)

Data packets travel across the Internet from router to router until they reach their final destination. At each router, packets arrive through incoming links and are transmitted through the appropriate outgoing link that leads to the next stop – which can be a router or the receiving end host – on their path to their ultimate destination.

If packets arrive for transmission over an outgoing link while another packet is being transmitted across that link, they are stored in a queue (or "buffer") for that link until it is their turn to be transmitted.[124] If packets destined for a specific outgoing link arrive faster than they can be

(continued)

example, if the service offers a constant performance that is specified in absolute terms and the specified performance is worse than the performance experienced by the best-effort service if the network is not congested. Huston (2012), p.2.

[120] While most network-neutrality-related discussions focus on services whose performance is better than best-effort service, a network that offers Quality of Service may also offer services that are worse than best effort service during times of congestion. For example, a class of service may provide a "less-than-best-effort" service ("scavenger class") that sends almost no traffic during times of congestion. See Shalunov & Teitelbaum (2001); Internet2 (Undated); Cisco Systems (2005b), Section "QoS Requirements of the Scavenger Class."

[121] How the performance of the service compares with best-effort service in the absence of congestion depends on the type of service. See footnotes 118 and 119 above and accompanying text.

[122] Even in an uncongested network, applications will still experience delay and may encounter losses. Data packets need to travel across the network, which takes time, and packets may get lost for reasons other than congestion. See Kurose & Ross (2010), pp. 36-40 (describing the different types of delay contributing to a packet's total end-to-end delay).

[123] The definition of congestion used throughout this section is derived from the definition of congestion in queuing theory. See *Box 8: Definitions of Congestion and Benefits from Quality of Service* below.

[124] On the following, see generally Kurose & Ross (2010), pp. 337-340, 653-654; Peterson & Davie (2012), 16-17, 479-480, 492-493. See also Reed (2009), paras 6-18. The text describes the scheduling and drop policy – FIFO queuing with tail drop – that is commonly used in the public Internet at the time of this writing. Each outgoing link has one queue. The router transmits packets over the link in the order in which they arrive ("first in, first out"). If the queue is full when a packet arrives, that packet is discarded ("tail drop"). Under some drop policies, the router may discards packets from its queue to make space for the arriving packet. Peterson & Davie (2012), pp. 492-493.

transmitted over that link, the number of packets in the queue increases. This may happen, for example, at routers that connect faster incoming links with slower outgoing links, or when different data transfers across the same link coincide.[125] As the number of packets in the queue increases, packets arriving for transmission across that link have to wait longer until they are transmitted, which increases the delay they experience. If the queue is full and cannot accommodate additional packets, the router starts dropping arriving packets, creating packet loss.

The end-to-end delay (or "latency") experienced by a packet indicates how long it takes the packet to travel from its origin to its destination. A packet's end-to-end delay consists of a number of components: how long it takes for the packet to be processed by the various routers along its path, how much time the packet spends in router queues waiting to be transmitted, how long the various routers need to transmit the packets onto the appropriate outgoing link and how long the packet needs to travel along the links from one router to the next.[126] The longer a packet has to wait in one or more router queues along its path, the higher its end-to-end delay.

Now consider an application that sends a number of data packets from one end host to another that travel along the same path ("data flow"). If the different packets spend varying amounts of time in router queues along their way, their end-to-end delay will vary. This variation in end-to-end delay is called jitter.[127] If all packets in a data flow have a similar end-to-end delay (e.g., because they all experience no queuing delay, or because all experience a similar, higher queuing delay), jitter is low. By contrast, if the end-to-end delay experienced by packets in the flow is highly variable (e.g., because some packets experience a lot of delay, while others experience little delay), jitter is high.

[125] See, e.g., CACM Staff (2012), p. 43; Peterson & Davie (2012), 16-17. For a more detailed description of potential reasons for congestion, see also *Box 9: Causes of Congestion in a Network with Low Average Utilization* below.

[126] More technically, a packet's end-to-end delay consists of the sum of all the processing delays, queuing delays, transmission delays and propagation delays the packet experiences along its path. Kurose & Ross (2010), pp. 37-40; Peterson & Davie (2012), pp. 46-47.

[127] Peterson & Davie (2012), pp. 54-55.

Box 8

Definitions of Congestion and Benefits from Quality of Service

Throughout this section, "congestion" denotes the building up of a queue for an outgoing link at a router, which may increase delay, jitter or packet loss (see *Box 7: The Relationship between Congestion, Delay, Jitter and Loss* above). This definition is derived from the definition of congestion used in queuing theory.[128] As explained in the text, Quality of Service only provides an improvement over best-effort service if this type of congestion exists.

By contrast, under a definition often used by network providers, congestion occurs if the average utilization of a link over a certain time period exceeds a certain threshold.[129]

While Quality of Service is useless in a network that never experiences congestion under the definition used throughout this section, is may still be useful in a network that is not congested under the definition used by network providers. Even in a network with low average utilization queues will build up occasionally.[130] Thus, a network that is not congested under the definition used by network providers may experience congestion under the definition used throughout this section and may therefore benefit from Quality of Service. As a result, the statement "Quality of Service is only useful if there is congestion" is correct only under this section's definition of congestion, but is false if the term congestion is used according to the network providers' definition.

In a network where average utilization is high, congestion will occur often and for extended periods of time. During periods of extended congestion, QoS-sensitive applications may become effectively unusable with best-effort service and may require a different class of service to function satisfactorily.[131] In such a network, users may find Quality of Service very valuable and may be very willing to pay for it.[132]

Adding capacity to reduce average utilization will reduce the amount of congestion. If average utilization is low, congestion will tend to occur less often and may cause less loss or

[128] See Bauer, Clark & Lehr (2009), pp. 8-9 and p. 10 ("According to [the queuing theory] definition, as soon as a queue starts to build traffic congestion is occurring.") See also Reed (2009), paras 7 and 9 ("In the Internet context, congestion manifests itself in routers or switches that forward Internet datagrams along the path between a particular source or destination. […] Congestion […] occurs when the amount of data that must travel through a particular link out of a particular router exceeds the data rate of that link for a long enough period such that a queue builds up.")

[129] See Bauer, Clark & Lehr (2009), pp. 10-11 and Reed (2009), paras 24-25. A network can be congested under the queuing theory definition, but not under the network provider's definition, and vice versa. Bauer, Clark & Lehr (2009), pp. 11. For a detailed analysis and comparison of different definitions of congestion, see ibid., pp. 8-13.

[130] See *Box 9: Causes of Congestion in a Network with Low Average Utilization* below and footnote 133 below.

[131] While Quality of Service may improve the performance of applications at a given capacity relative to a single best-effort service operating over the same capacity, even Quality of Service needs a certain amount of capacity to provide satisfactory performance. See, e.g., Kurose & Ross (2010), pp. 664-665; Peterson & Davie (2012), p. 553.

[132] As discussed above, proponents of a ban on Quality of Service are concerned that network providers may have an incentive to operate networks in this state since this increases users' willingness to pay for Quality of Service. While this incentive exists, it can be constrained in other ways that fall short of banning all forms of Quality of Service. See footnotes 104 to 106 above and accompanying text.

delay. But even a network with low average utilization will experience occasional congestion.[133] For a number of reasons, queues will form temporarily even when average utilization is low, and if the resulting increase in delay, jitter or loss exceeds the amount that a QoS-sensitive application can compensate for,[134] the performance of that application will suffer. (See *Box 9: Causes of Congestion in a Network with Low Average Utilization* below.)

Box 9

Causes of Congestion in a Network with Low Average Utilization

Congestion will occur even in a network with low average utilization. For a number of reasons, queues will form temporarily, creating congestion, even when average network utilization is low.

Many Internet applications are bursty: their peak rate is much higher than their average rate.[135] Under these circumstances, focusing only on average utilization is misleading. The capacity of the links along a bursty application's path may be more than sufficient to transmit data at that application's average rate without delay. But if the application's peak rate is higher than a link's available capacity, the application will temporarily send data faster than the link can transmit, filling up the link's queue until the burst subsides. More generally, whether a specific link gets congested at a specific point in time depends on whether the actual data rates of the various applications sharing the link at that moment exceed the link's capacity, not on the average data rates of these applications.

On today's Internet, bursty applications create challenges for interactive applications. For example, applications such as web browsing or streaming video send short bursts of data packets that may temporarily fill queues; when the burst ends, the queues drain quickly. This not only increases the delay experienced by other applications that are transferring data over the same link at the same time; the rapid building up and emptying of queues also increases jitter. The increase in jitter and

(continued)

[133] This insight is well established in the literature. See, e.g., Carpenter & Nichols (2002), pp. 1482-1483 ("It is not necessary for the network's long-term utilization to be high for this to occur; the traffic burstiness mentioned above can lead to congestive incidents even when average utilization is modest", p. 1482); Bauer, Clark & Lehr (2009), pp. 6, 11, 16, 32 ("However, because demand is not smooth and fluctuates stochastically over time at many different time-scales and because the available capacity the Internet varies across the network, congestion events may arise commonly even in network that may be considered to be generally "overprovisioned", p. 6); Clark, Lehr & Bauer (2010), p. 10.

[134] The ability of an application to compensate for increases for delay, jitter or loss resulting from congestion at the end host is systematically limited. Whether and to what extent a certain application will be able to compensate, depends on the performance requirements of that application and the size of the congestion-related increase in delay, jitter or loss. For example, an application that streams video from a server to the user (such as YouTube or Netflix) can tolerate or compensate for a higher level of jitter than an interactive, real-time video-conferencing application. Applications compensate for jitter by buffering data for playback. Compensating for higher jitter requires a larger buffer. The resulting increase in delay will be more tolerable for streaming stored video than for interactive real-time applications. See, e.g., Peterson & Davie (2012), pp. 532-534; Järvinen, et al. (2012), p. 2.

[135] The relationships described in the text are explained in more detail in, e.g., Peterson & Davie (2012), pp. 54, 540; Kurose & Ross (2010), pp. 40-42; Reed (2009), paras 16-18.

delay harms applications such as interactive voice and video applications or online gaming applications which need low jitter or delay.[136] Recent changes to transport protocols[137] and operating systems[138] have increased the amount of data a single TCP connection may send, which increases the potential peak rate at which bursts may occur. In addition, today's browsers transmit data over several parallel transport layer connections simultaneously, creating even larger bursts of data that can easily fill up a link's queue.[139]

Applications that upload or download a lot of data using the Transmission Control Protocol (TCP) (e.g., for uploading a video to YouTube, sending or receiving emails with large attachments or backing up data to the cloud) pose challenges of a different kind. They create long-lived data flows that cause standing queues in routers for the duration of the flow, which increases delay for other applications trying to transfer data at the same time.[140]

Moreover, TCP is designed to increase its transmission rate until it uses all available bandwidth and to reduce its transmission rate when it detects congestion. Thus, as long as the amount of data to be sent by an application is sufficiently large, TCP by design creates instantaneous congestion, even in a well-provisioned network.[141]

While many users may be willing to tolerate the temporary lower performance associated with occasional congestion, some users may value more reliable performance. Many users use Skype even though the quality of the call often varies over the duration of the call and calls break up occasionally. While Skype's quality will often be good enough for them, at least some of these users (or users who are not using Skype in the current Internet because Skype's performance is not good enough for them) may value (and be willing to pay for) the option of

[136] See Järvinen, et al. (2012) (discussing the problem and presenting the results of an experiment that demonstrates the problem).

[137] For example, a recent IETF experimental standard proposes an increase in the permitted upper bound for TCP's initial window (IW) to 10 segments depending on the maximum segment size (Chu, et al. (2013)). These changes allow each new TCP connection "to send as much as 2.5 times data as in the past."(Gettys (2011), p. 3). For an analysis of the impact of these changes on other applications that are transferring data at the same time, see Gettys (2011); Chu, et al. (2013), pp. 5, 9-10, 13-14; Järvinen, et al. (2012).

[138] For example, Windows XP did not implement TCP window scaling, and therefore, the amount of packets it sent before it received an acknowledgment was limited. As a result, Windows XP was less likely to saturate links than newer versions of the Windows operating system, which do implement TCP window scaling. At the time of this writing, Mac OS X, Linux and Windows operating systems after XP all implement window scaling. See, e.g., Gettys & Nichols (2012), p. 64; CACM Staff (2012), p. 44-45. TCP windows scaling is a TCP option that makes it possible to increase the size of TCP's receive window beyond 65K bytes, the maximum size of the receive window under normal TCP (Jacobson, Braden & Borman (1992)).

[139] For example, current browsers open six or more (e.g., fifteen) TCP connections to a single website. In addition, some websites ("sharded websites") are engineered to appear as if data is coming from different domains, which tricks the browser into allowing even more TCP connections. See, e.g., Gettys (2011), pp. 3-4; Järvinen, et al. (2012), p. 1.

[140] See, e.g., Nichols (2012), pp. 4-5; Järvinen, et al. (2012), p. 2. This problem has been exacerbated by large buffers ("bufferbloat") in the access networks and elsewhere. The larger the buffer, the longer the queue can become, and the longer the delay experienced by packets that are arriving when the queue is almost full. See Gettys & Nichols (2012); CACM Staff (2012).

[141] See, e.g., Bauer, Clark & Lehr (2009), p. 16; Clark, Lehr & Bauer (2010), p. 10.

using a different class of service that would allow them to get reliably good or even excellent call quality for selected Skype calls. Hearing-impaired users that rely on sign-language to communicate may value perfect picture quality in video telephony more than "normal" users. A traveler on a business trip may be willing to tolerate occasional glitches and break-ups in the video chat when saying good-night to her children at home, but may need high-quality, predictable performance when using the same application to give a talk at a conference.[142] Thus, the absence of classes of service that provide more reliable (or potentially better) performance than best-effort service may hurt users who would value being able to take advantage of them when needed.

In addition to giving users the option to improve the performance of existing applications, Quality of Service may enable the development of new applications that cannot function in today's public Internet since they have requirements that a best-effort network cannot support. For example, a best-effort network cannot provide any guarantees with respect to throughput, jitter or delay, making it impossible to support applications that strictly need guaranteed throughput, jitter or delay.[143] More generally, there may be applications that may benefit from the availability of services other than best-effort. Thus, it is at least possible that a total ban on Quality of Service may reduce innovation in QoS-sensitive applications, harming users who would have benefited from these applications.[144] In conversations, proponents of a ban on Quality of Service often reject this argument as hypothetical. They would like to see compelling examples of applications that require Quality of Service before they are willing to consider the possibility that Quality of Service may foster application innovation.[145] Economic theory and the history of general purpose technologies suggest, however, that it is usually not possible to predict in advance how a general purpose technology will be used and which potential uses will be successful.[146] Throughout the history of the Internet, most Internet applications that later became highly successful either were not envisaged by the designers of the network or were met by widespread skepticism when they first became available. This was true, for example, for e-mail, the World Wide Web, eBay or search engines.[147] Thus, that we cannot imagine socially beneficial applications that require Quality of Service does not mean that such applications do not exist. Instead, the history of the Internet suggests that when a large, diverse group of

[142] The interest in getting more reliable performance may vary across users and, for the same user, over time, and any rules for Quality of Service should reflect that. How this insight affects which forms of Quality of Service a non-discrimination rule should allow is discussed in footnotes 424 to 430 below (discussing forms of Quality of Service that treat like traffic alike) and accompanying text and footnotes 481 to 482 below and accompanying text (discussing certain forms of user-controlled Quality of Service).

[143] It is not possible to construct a transport layer service that guarantees delay (or bandwidth), if, as in the current best-efforts Internet, the Internet layer does not guarantee delay (or bandwidth). van Schewick (2010d), p. 142, Box 5.3; Kurose & Ross (2010), pp. 97, 201.

[144] Litan & Singer (2007), pp. 569-570.

[145] This argument has come up repeatedly in personal discussions with proponents of a ban.

[146] See, e.g., Rosenberg (1996). The next few sentences draw on van Schewick (2010d), pp. 301-302.

[147] For a more detailed discussion of these examples, see van Schewick (2010d), pp. 301-304.

innovators is allowed to innovate under the right conditions,[148] they will find ways to use the Internet's functionality that those who originally designed that functionality had not necessarily thought of, and at least some of the resulting applications or uses will create significant social value.[149]

Finally, in situations where a user's desire for bandwidth exceeds the amount of bandwidth available to him (for example, because the size of the access link is limited or the network provider limits the amount of bandwidth available to individual subscribers during peak times when average network utilization is high), allowing certain forms of Quality of Service may enable users to use that limited amount of bandwidth more efficiently.[150]

Network providers could reduce the likelihood of congestion even further by increasing capacity so that "the capacity of individual links is significantly larger than the peak average traffic of all users."[151] This solution is called "over-provisioning."[152] Provisioning links significantly above the peak average traffic of all users of the link requires considerably more capacity (and will be considerably more expensive) than ensuring low average utilization. For example, in 2006, representatives of the research network Internet2 suggested that over-provisioning residential access networks or, as they described it, providing "the abundance of bandwidth [that] ensure[s] that the odds of network congestion are minimized" would require offering a one gigabit per second connection to residential users (one gigabit per second equals 1,000 megabit per second).[153] Since then, the demands and capabilities of end devices and applications have evolved rapidly, so the capacity required to over-provision access networks today will most likely be higher. For example, a single TCP connection on a PC can send data at a rate of hundreds of megabits per second,[154] so a single user could easily create peak rates of more than a gigabit per second by opening several TCP connections simultaneously. Moreover, TCP is designed to

[148] These conditions include the factors described above: innovation without permission, innovation without fear (or application-agnosticism) and user choice. See footnotes 50 to 59 above and accompanying text. If we want Quality of Service to foster application innovation, any rule allowing Quality of Service must ensure that it is offered in a way that preserves these factors. To see how this insight affects which forms of Quality of Service a non-discrimination rule should allow, see footnotes 403 to 434 and accompanying text below (discussing forms of Quality of Service that treat like traffic alike) and footnotes 479 to 483 and accompanying text below (discussing certain user-controlled forms of Quality of Service).

[149] For a more detailed discussion of this argument based on economic theory and the history of specific Internet applications, see van Schewick (2010d), chapter 8.

[150] See the discussion of application-agnostic network management coupled with user-controlled differentiation in Section "Ban application-specific discrimination, allow application-agnostic discrimination" below.

[151] Reed (2009), para 23.

[152] As Bauer, Clark & Lehr (2009) explain, a common approach to provisioning , which provisions for expected peak demand over some time period, may often result in networks that are over-provisioned over considerable amounts of time: "Indeed, a common approach to managing resource sharing is to provision for expected peak demand over some time period, and because many network investments need to be made in relatively large fixed increments and over an investment time horizon that takes months or more, capacity is provisioned in advance of realized demand. Thus, during offpeak periods (which may be measured in periods of hours or days) and over the life of infrastructure investments (which may be measured in periods of months or years), there may be significant amounts of time when the network is over-provisioned relative to offered demand." (ibd., p. 6)

[153] Corbató & Teitelbaum (2006), pp. 2-3; Bachula (2006), p. 3; Kuhns (2006).

[154] Bauer, Clark & Lehr (2009), p. 16.

use all available bandwidth. As long as it has data to send, TCP speeds up until it detects congestion, so any network over which TCP is used will always experience some temporary congestion.[155] Finally, even in an over-provisioned network, data may travel from faster to slower links, coinciding data transfers may temporarily exceed the capacity of a link, or unexpected spikes in demand may exhaust a link's capacity, all of which creates congestion as well.[156] Thus, while over-provisioning will further reduce the probability of congestion, it cannot eliminate it.[157] Due to the low likelihood of congestion, a network that is truly over-provisioned will probably be able to support most QoS-sensitive applications most of the time. But even in such a network, Quality of Service may still be useful as "insurance" against the residual risk of congestion.[158]

In sum, the value of Quality of Service is not restricted to networks with high average utilization which are often congested. While Quality of Service is only useful if there is congestion (i.e. if queues build up in routers), increasing capacity does not necessarily prevent congestion and therefore, Quality of Service may be useful in networks with more capacity as well. In networks that have low average utilization, but are not over-provisioned,[159] Quality of Service may give users the option to improve the performance of existing applications by using classes of service that provide more reliable or potentially better performance than best-effort service if congestion occurs. Quality of Service may also enable new applications that we have not thought of yet that cannot function in a best-effort Internet or that would benefit from classes of service other than best-effort. And it may allow users whose bandwidth is limited to use that limited amount of bandwidth more efficiently. While the relative value of Quality of Service is likely to decline as a network's capacity approaches the capacity required for over-provisioning, Quality of Service may provide benefits even in over-provisioned networks by allowing users to

[155] For a detailed explanation of this point, see Bauer, Clark & Lehr (2009), p. 16. See also Clark, Lehr & Bauer (2010), p. 10 ("Some observers seem to argue that a preferred alternative to adding QoS is simply to expand capacity, or equivalently, over-provisioning of the network so that congestion does not occur. […] We believe that this line of reasoning is flawed […]. Since TCP tries to go as fast as possible unless it is being artificially throttled (as does occur today in some cases), congestion will occur somewhere along the path, if only in the server itself.")

[156] See also *Box 9: Causes of Congestion in a Network with Low Average Utilization* above.

[157] Bauer, Clark & Lehr (2009), p. 5 (For a number of reasons [BvS: see footnote 152 above], "there may be significant amounts of time when the network is over-provisioned relative to offered demand. During such periods, the network may appear to be relatively uncongested. However, because demand is not smooth and fluctuates stochastically over time at many different time-scales and because the available capacity the Internet varies across the network, congestion events may arise commonly even in a network that may be considered to be generally "over-provisioned.") See also Teitelbaum & Shalunov (2002), Section "QoS and DoS" ("Although well-provisioned networks deliver very good typical performance, they will, in general, deliver unpredictable service and, in the worst case, no service.")

[158] Teitelbaum & Shalunov (2002), Section "Premium Service" (arguing that Premium service would be valuable even in an over-provisioned network like Internet2: "Premium service is about guaranteeing service quality. In essence, it is about removing a component of unreliability from the system--the probability that a network transaction fails because of network congestion. Although typical performance may be perfect, there would be considerable value in being able to assure that important sessions receive perfect network performance."); Teitelbaum & Shalunov (2003), p. 149 (arguing that the ultimate goal of Quality of Service in a well-provisioned network is "to eliminate or bound the risk that preferred traffic will experience congestion", and that this function is valuable even in a well-provisioned network).

[159] Over-provisioning requires considerably more capacity than ensuring low average utilization, so a lot of networks may belong to this category.

protect selected applications against the residual risk of congestion. Thus, banning Quality of Service has social costs, and these costs exist over a wide range of network capacities.

While some proponents of banning all forms of Quality of Service argue that the costs of a ban are negligible since the needs of QoS-sensitive applications can be met by increasing capacity, some supporters of a ban make a stronger claim: According to them, banning Quality of Service does not have social costs because over-provisioning is economically and technologically more efficient than offering Quality of Service, so banning Quality of Service only prohibits a technical solution that is less efficient anyways.[160] Quality of Service makes the network more complex and is more difficult to manage than a single best-effort service. Network engineers have debated for years whether the benefits of Quality of Service outweigh the added complexity and cost, or whether over-provisioning is more efficient.[161] After developing and successfully testing Quality of Service technology in the research network Internet2 for several years, Internet2 researchers suspended the effort indefinitely. While they acknowledged that being able to protect important applications against the risk of congestion is valuable even in an over-provisioned network, they concluded that "the costs […] would be too high relative to the perceived benefits" and that over-provisioning was the more efficient solution.[162] In congressional testimony and elsewhere, representatives of Internet2 have used this experience to argue in favor of network neutrality rules that ban Quality of Service.[163]

While introducing Quality of Service creates costs, over-provisioning – which requires considerably more capacity than ensuring low average utilization – is not costless, either. Routers' processing power, the administrative costs of deploying and managing Quality of Service technology and the costs of deploying additional capacity may differ across different types of networks and may change over time. For example, backbones may be easier to over-provision than access networks because they can take advantage of statistical aggregation.

[160] See, e.g., Open Internet Coalition (2010), pp. 43-46 (citing the experience of Internet2 as support for the proposition that "[t]he most technologically and economically efficient means of managing Internet traffic is by increasing capacity." Ibid., p. 43).

[161] Kurose & Ross (2010), pp. 602-604; Bell (2003), pp. 116-119 (discussing the trade-off in the context of enterprise networks); Davie (2003), pp. 134-135. In proponents of over-provisioning, see Odlyzko (1998); Odlyzko (1999); Bricklin (2003); Huston (2012) and the sources cited in footnote 163 below.

[162] Teitelbaum & Shalunov (2002), Section "Abstract" ("The costs of Premium are too high relative to the perceived benefits.") and Section "Premium Service" ("Premium service on a well-provisioned network would do little to change packet forwarding under normal conditions. Internet2 networks are generally well-provisioned and almost always lightly loaded. Packet loss and jitter experienced by best-effort traffic on Internet2 paths is almost always zero or is due to non-congestive causes. […] Nevertheless, […] Premium service is about *guaranteeing* service quality. In essence, it is about removing a component of unreliability from the system—the probability that a network transaction fails because of network congestion. Although typical performance may be perfect, there would be considerable value in being able to assure that important sessions receive perfect network performance.") While the document discusses the QBone Premium service, an interdomain virtual leased-line IP service built on diff-serv forwarding primitives, the authors claim that the reasons for suspending the deployment of the QBone Premium service "apply not just to Premium, but to any IP quality of service (QoS) architecture offering a service guarantee." (ibd., Section "Abstract")

[163] Bachula (2006), pp. 2-3; Kuhns (2006); Corbató & Teitelbaum (2006); Teitelbaum & Shalunov (2002). The Internet2 experience is often cited by proponents of a ban on Quality of Service. See, e.g., Open Internet Coalition (2010), pp. 43-45 (citing the experience of Internet2 as support for the proposition that "[t]he most technologically and economically efficient means of managing Internet traffic is by increasing capacity." Ibid., p. 43).

Over-provisioning research networks whose users are already attached to high-speed campus networks may be less costly than over-provisioning residential access networks. The complexity and costs of deploying and running Quality of Service may be lower in enterprise networks, where the same entity controls all parts of the network infrastructure (including the end hosts), than in multi-provider networks.[164] Today, many corporate intranets use Quality of Service; large Internet service providers give business customers the option of buying different classes of service.[165] Thus, whether over-provisioning is more efficient than introducing Quality of Service may differ depending on the circumstances and may change over time.

The debate over the relative costs and benefits of over-provisioning and Quality of Service is an important debate worth having. But whatever the merits of this debate from a technical perspective, arguments over the relative cost-efficiency of alternative technical solutions should be irrelevant for the regulatory debate over network neutrality rules.

Network neutrality rules are the result of a trade-off.[166] They impose some constraints on the evolution of the network in order to allow the Internet to continue to foster application innovation, preserve user choice or foster democratic discourse. Policy makers need to decide whether restrictions on the evolution of the network (here: banning Quality of Service) are necessary to protect the values that network neutrality rules are designed to protect. If the restrictions are not necessary to protect these values, they should not be imposed. By contrast, whether introducing Quality of Service makes sense from a technical or business perspective is a question that should be left to network engineers and network providers.[167] If regulators adopt non-discrimination rules that allow certain forms of Quality of Service, they do not pick winners and losers in this debate. Such non-discrimination rules do not require network providers to introduce Quality of Service; they only allow them to do so within the constraints imposed by the rules. If network providers decide that over-provisioning offers a better cost-benefit trade-off than offering Quality of Service in line with the rules, they are free to go down that route.

In sum, while there are legitimate concerns about the consequences of allowing Quality of Service on competition among applications or investment in the network, these concerns can be mitigated without totally banning Quality of Service. Different forms of Quality of Service have different social benefits and social costs, so a more nuanced treatment is needed. While the

[164] In addition, in enterprise deployments, the entity that incurs the costs of deploying and running Quality of Service also reaps the benefits. By contrast, the business model (and therefore, the expected benefits) associated with introducing Quality of Service in the public multi-provider Internet may be less clear, which makes it more difficult to justify the high costs of operating Quality of Service across the networks of multiple providers. See Davie (2003); Huston (2012), p. 4.

[165] On enterprise deployments, see Peterson & Davie (2012), p. 554; Davie (2003), pp. 133-134. On QoS-offerings for business customers, see, e.g., Verizon (2006); von Bornstaedt (2012).

[166] See footnotes 40 to 45 above and accompanying text.

[167] Of course, the constraints imposed by a non-discrimination rule that allows all or some forms of Quality of Service will influence network providers' private costs and benefits of over-provisioning and Quality of Service. For example, other things being equal, introducing Quality of Service may be more attractive under a regime that allows network providers to charge whomever they like for the provision of different classes of service, and less attractive under a regime that prohibits network providers from charging for Quality of Service.

value of Quality of Service may decline as network capacity increases, Quality of Service may be useful over a wide range of network capacities, not just in networks with high average utilization. In networks that have low average utilization without being over-provisioned, Quality of Service may allow users to improve the performance of existing applications, may enable new applications that benefit from the availability of different classes of service, and may enable users whose bandwidth is limited to use that bandwidth more efficiently. Ensuring low average utilization requires considerably less capacity than over-provisioning, so many networks may belong to the category just described. In an over-provisioned network, Quality of Service offers users the option of protecting applications against the residual risk of congestion. Thus, at least some forms of Quality of Service may provide social benefits over a wide range of network capacity. At the same time, the social costs of offering Quality of Service can be limited through appropriate rules. Under these circumstances, requiring network providers to treat every packet the same would be too restrictive, constraining the evolution of the network more than absolutely necessary to protect the values that network neutrality is designed to protect.

All-or-nothing approaches are appealing because they impose clear obligations that are easy to enforce. All industry participants know what to expect and can adjust their behavior accordingly. These advantages come at a cost. All-or-nothing approaches treat all forms of discrimination in the same way. As a result, they are either over-inclusive, banning forms of discrimination that would be socially beneficial, or under-inclusive, allowing socially harmful forms of discrimination.

CASE-BY-CASE APPROACHES

A second set of approaches wants regulators to determine case-by-case whether discriminatory behavior that falls short of blocking should be forbidden. Proponents of these approaches recognize that some forms of differential treatment will be socially harmful, while others will be socially beneficial. As a result, they reject a blanket ban on discrimination as over-inclusive. At the same time, they doubt that it is possible to distinguish socially beneficial from socially harmful differential treatment in advance. According to them, this determination is best made ex post, when the facts that will allow an accurate assessment of the practice, such as motivations for and impact of the practice, are known.[168] To support their proposal, they point to the example of antitrust law, which evaluates behavior that may be anticompetitive or pro-competitive depending on the circumstances after the fact on a case-by-case basis.[169]

Approaches in this group differ along two dimensions:[170] the degree to which they prescribe the standard which regulators should use to assess specific discriminatory behavior,

[168] See, e.g., Weiser (2003), pp. 75-76; Yoo (2007), pp. 515-517; Hahn, Litan & Singer (2010), p. 368.

[169] Weiser (2003), pp. 75-76; Yoo (2007), pp. 515-516.

[170] Proposals also differ on which institutional actor should perform the adjudication (in the US, the FTC (see, e.g., Nuechterlein (2009), pp. 57-65) or the FCC (see, e.g., Weiser (2003), pp. 75, 77-78; Hahn, Litan & Singer (2010), p. 374 (listing reasons for enforcement by the FCC without taking a side in the debate)). This question is outside the

and to what extent they are able to capture the instances of discrimination about which network neutrality proponents are concerned. Taken together, these two characteristics determine how likely it is that an actor that encounters discrimination that network neutrality proponents would classify as harmful will prevail in the future.

Approaches at one end of the spectrum specify the standard for separating socially harmful from socially beneficial discrimination, but the standard would not capture many instances of discrimination that network neutrality proponents are concerned about, classifying them as socially beneficial. Thus, these approaches would often make it impossible to successfully challenge behavior that network neutrality proponents would view as harmful. Proposals that suggest using an antitrust framework are an example of this type of approach.

Approaches at the other end of the spectrum do not specify the standard at all. As a result, the proposed rule is consistent with interpretations that capture all relevant (from the perspective of network neutrality proponents) instances of discrimination and with interpretations that do not. Thus, under such a rule it is at least possible, but not certain, that a challenge to behavior that network neutrality proponents deem harmful will be successful.[171] The draft Open Internet Rules circulated by the FCC Chairman in early December 2010 are an example of this type of approach. They banned "unreasonable discrimination," without specifying how this term should be interpreted.

Ban Discrimination That Violates an Antitrust Framework

The first set of proposals in this group suggests using an antitrust framework to distinguish socially beneficial from socially harmful discrimination.[172] These proposals interpret the concerns raised by proponents of network neutrality regulation as concerns about anticompetitive vertical leveraging or vertical foreclosure[173] and apply the framework used to evaluate vertical leveraging and vertical foreclosure claims under the US antitrust laws to determine whether discriminatory conduct should be banned.[174] The term "vertical leveraging" describes a situation in which a firm that has a monopoly in one market – here, a provider of

scope of this paper. The FTC is only a plausible option for those who base network neutrality regulation on an antitrust framework. If network neutrality regulation is based on a broader framework, as assumed in this paper, the FCC is the right agency to enforce the rules, since it is used to and tasked with applying a broader public interest standard.

[171] Currently, there is no approach that clearly defines the standard in a way that would capture all relevant instances of discrimination. The final Open Internet Rules adopted by the Federal Communications Commission in December 2010 provide more detail on how to distinguish socially beneficial from socially harmful behavior, but they still leave a considerable amount of uncertainty.

[172] See, e.g., Weiser (2003), pp. 74-84; Nuechterlein (2009), pp. 20-65; Yoo (2007), pp. 508-517; Yoo (2008), pp. 245-247, 257-261; Yoo (2009), pp. 81-83; Becker, Carlton & Sider (2010), p. 508; Hahn, Litan & Singer (2010), pp. 374-379; Sidak & Teece (2010), pp. 562-563; Hazlett & Wright (2011), Section V., pp. 31-32; Ohlhausen (2012), pp. 10-17.

[173] See, e.g., Weiser (2003), pp. 71-74; Nuechterlein (2009), p. 35; Becker, Carlton & Sider (2010), pp. 501-502, p. 508; Hazlett & Wright (2011), Section V., pp. 31-40; Ohlhausen (2012), pp. 10-17.

[174] Proposals differ both in the level of detail with which they describe the framework and in the exact criteria they use to distinguish socially harmful from socially beneficial discrimination. The text focuses on what seem to be the unifying threads in the various proposals.

Internet access services – "abuses" or "leverages" its market power in the first market to obtain an unfair[175] advantage in a second, vertically related market – for example, in the market for a specific application.[176] The term "vertical foreclosure" applies to situations in which a monopolist in a primary market – i.e., a provider of Internet access service – uses its market power in the first market to deny firms in a second, vertically related market – i.e. the market for a specific application – access to that second market.[177] Over the years, the views of US antitrust scholars and courts towards these practices have evolved considerably. Today, US antitrust law condemns vertical leveraging or vertical foreclosure only if the exclusionary conduct meets the criteria of Section 2 of the Sherman Act, which prohibits monopolization or attempts to monopolize.[178]

This standard does not capture all instances of discrimination that threaten the values that network neutrality rules are designed to protect. Challenges to discriminatory behavior that network neutrality proponents deem socially harmful may fail for one of four reasons:

First, US antitrust law only condemns a network provider's discriminatory behavior that affects the market for a specific application, content, or service, if the network provider participates in that market or is affiliated with a participant in that market. As Areeda and Hovenkamp's antitrust treatise explains, "[e]ven the most expansive formulations of 'leveraging' […] limit the concept to situations where the defendant [i.e. the primary good monopolist] actually does or intends to do business in the secondary market. Mere injury to firms in a vertically related market in which the defendant does not operate cannot be leveraging, for nothing is being leveraged."[179,180]

By contrast, network neutrality proponents are also concerned about discrimination in application markets in which the network provider does not participate. For example, network providers may have an incentive to block unwanted content that threatens the company's

[175] Hovenkamp (2011b), §7.9, p. 348.

[176] Areeda & Hovenkamp (2010), p. 130 para 652 a.

[177] Hovenkamp (2011b), §10.6b2, pp. 462-463.

[178] Hovenkamp (2011b), §7.9, pp. 348-349. Areeda & Hovenkamp (2010), para 652b, pp. 134-140. Tying and exclusive dealing are evaluated according to different criteria, but most of the behavior that network neutrality proponents are concerned about does not qualify as tying or exclusive dealing.

[179] Areeda & Hovenkamp (2010), para 652, p. 139 (discussing the question in the context of monopoly leveraging claims under Section 2 of the Sherman Act). Behavior by a monopolist that negatively affects competition in a complementary market in which the monopolist does not operate does not violate Section 5 FTC Act, either. The Federal Trade Commission (FTC) originally adopted a different view in its Order against Official Airline Guides, but this decision was reversed by Official Airline Guides v. Federal Trade Commission 630 F.2d 920 (2d Cir. 1980) (United States Court of Appeals (1980)). Since then, FTC has signaled that it does not deem this type of behavior actionable under Section 5 FTC Act any more. See Federal Trade Commission (2003).

[180] In the network neutrality context, vertical integration by network providers into applications is viewed as a prerequisite for regulatory intervention by, e.g., Sidak & Teece (2010), p. 563 (criticizing the non-discrimination rule proposed in the FCC's Open Internet NPRM as overbroad, because it applies to network providers regardless of whether they are "vertically integrated into providing competing content."); Hahn, Litan & Singer (2010) , pp. 373, 375-276 ("[I]n the absence of vertical integration into the content space, a BSP [Broadband Service Provider] will lack any incentive to discriminate between content providers who demand the same service.", p. 373); Cave & Crocioni (2011), p. 65 (explaining that consumer harm from exclusion depends on ISPs being vertically integrated into applications).

interests or does not comply with the network provider's chosen content policy. This incentive is independent of whether the network provider operates in the market for the affected content. In the examples of content-based discrimination that are often mentioned in the debate (e.g. Telus/Voices for Change; Verizon Wireless/NARAL Pro Choice, Apple/iSinglePayer (see *Box 10: Examples of Content-Based Discrimination* below)), none of the content providers whose content was blocked was competing with the network provider. Similarly, a network provider may have an incentive to exclude or slow down selected bandwidth-intensive applications to manage bandwidth on its network, even if the network provider does not offer a competing application itself.[181] In these cases, the resulting harm – users' inability to participate in social, cultural or democratic discourse related to the blocked content, their inability to use the Internet in the way that is most valuable to them, or application developers' difficulty to obtain funding for an application – is caused by the discriminatory behavior as such and is independent of whether the network provider is active in the market or not.

Box 10
Examples of Content-Based Discrimination

> In 2005, Telus, Canada's second largest ISP, blocked access to a Web site that was run by a member of the Telecommunications Workers Union. At the time, Telus and the union were engaged in a contentious labor dispute, and the Web site allowed union members to discuss strategies during the strike. In 2007, Verizon Wireless rejected a request by NARAL Pro-Choice America, an abortion rights group, to let them send text messages over Verizon Wireless' network using a five-digit short code. In the same year, AT&T deleted words from a Webcast of a Pearl Jam concert in which the singer criticized George W. Bush. In 2009, Apple rejected an application called "iSinglePayer" that advocated for a single payer health insurance system as "politically charged." Verizon Wireless, AT&T and Apple all argued that the rejected or deleted content violated their content policies. They later changed their view after the incidents had been widely reported. While the latter three examples are not direct examples of ISPs restricting content on their networks (Verizon Wireless restricted a service on its wireless mobile network, not the wireless Internet; Apple acted as provider of the Apple App Store; and AT&T acted in its role as a content provider, not as ISP), it is easy to imagine virtually identical incidents in which an ISP enacts a content policy and restricts content on its network accordingly.[182]

Second, US antitrust law only condemns vertical leveraging or vertical foreclosure as monopolization or attempted monopolization under Section 2 of the Sherman Act, if they are reasonably capable of monopolizing the primary market or the secondary market.[183] Thus, to be classified as socially harmful under an antitrust framework, a network provider's discriminatory behavior in the market for a specific application must be reasonably capable of creating, increasing or maintaining monopoly power in the market for that application or in the market for

[181] See van Schewick (2010d), pp. 264-266.

[182] This paragraph is adopted from van Schewick (2009), p. 32. For a more detailed description of these examples, see van Schewick (2010d), pp. 266-269. On Apple/iSinglePayer, see LambdaJive (2009); Singel (2009).

[183] Areeda & Hovenkamp (2008), pp. 130-132, para 652a, pp. 140-142, para 652b.

Internet access services.[184] By contrast, network neutrality proponents may classify discriminatory behavior as socially harmful even if the behavior is unlikely to monopolize the application market or the market for Internet access services.

US antitrust law generally only condemns exclusionary conduct if it there is a reasonable likelihood that the behavior will harm competition, and not just competitors, by worsening the structure or performance of the affected market.[185] In the case of Section 2 of the Sherman Act, the behavior must be reasonably capable of creating, increasing or maintaining a monopoly, or of producing the higher prices or lower output or quality that attend monopoly. A firm's exclusionary behavior that just harms one or more competitors (e.g., by enlarging that firm's market share at the expense of its competitors) without creating or sufficiently threatening the higher prices or lower output or quality associated with monopoly is outside the scope of Section 2 of the Sherman Act.[186] Thus, to be condemned as socially harmful under an antitrust framework, a network provider's discriminatory conduct in the market for a specific application would have to drive affected applications from the market for that application, prevent new entry into an application market that the network provider has already monopolized, or impair the application provider's ability to compete effectively by forcing it to operate at a less efficient scale.

This requirement may be difficult to meet:[187] In many cases, the market for the application that is being discriminated against will be national in scope (see *Box 11: Relevant Geographic Market* below*),* while the network provider's customers only make up a part of the nation's Internet access customers. For example, in the US, the four largest broadband Internet access providers currently serve 25% (Comcast), 19% (AT&T), 14% (TimeWarner) and 11% (Verizon) of the nation's broadband Internet access customers.[188] Whether a network provider's discriminatory behavior will be capable of driving the application from the market or of preventing the application provider from reaching its minimum efficient scale in a way that unreasonably restrains the application's ability to compete effectively depends on the number of

[184] Sometimes, the discriminatory behavior may be designed to protect a network provider's existing monopoly in a third market, for example in the market for multi-channel video distribution or for telephony services. In this case, it is sufficient if the discriminatory behavior sufficiently threatens to perpetuate the network provider's monopoly in that market. It is unclear, however, whether the proponents of using an antitrust framework to evaluate complaints about non-discrimination would share this view. See, e.g., the discussion of the BitTorrent and Vonage cases by Hahn, Litan & Singer (2010).

[185] See Hovenkamp (2011a), p. 74, para 1802b ("The concern of the antitrust laws is with injury to 'competition,' which generally means injury resulting in lower output and higher prices in a properly defined market."). In the network neutrality context, a number of scholars explicitly evaluate discriminatory behavior based on whether it creates "harm to competition" in the antitrust sense. See, e.g., Farber, et al. (2007); Becker, Carlton & Sider (2010), pp. 501-502; Hahn, Litan & Singer (2010), p. 377.

[186] Areeda & Hovenkamp (2008), pp. 140-142, para 652c.

[187] See, e.g., Yoo (2005), pp. 71-73; Sidak (2006), p. 470, 472-473; Verizon & Verizon Wireless (2007), pp. 49-52; Sidak & Teece (2010), p. 566; Hahn, Litan & Singer (2010), pp. 371-372, 376; Litan & Singer (2007), pp. 556-557. This question is often discussed in the context of network providers' ability to foreclose applications. See, e.g., Hahn, Litan & Singer (2010), pp. 371-372.

[188] This data is current as of the end of the first quarter 2014. Market share calculations are based on Leichtman Research Group (2014).

foreclosed Internet access customers relative to the overall number of Internet access customers, on the size of any economies of scale in the market for the application, and on the size of the cost disadvantage associated with operating at a less than efficient scale. While many Internet applications are subject to significant economies of scale due to large fixed costs and low marginal costs or due to network effects,[189] exclusion from access to one ISP's customers may not create the type of anticompetitive harm that antitrust law is concerned about.[190] In such a case, an antitrust framework would not classify the exclusionary conduct as socially harmful.

Box 11
Relevant Geographic Market

> The market for an application is national in scope, if the application or content appeals to consumers nationwide instead of only to consumers in a particular locality. For example, local yellow pages for a specific region primarily appeal to consumers in that region. Thus, the relevant geographic market is local. By contrast, an Internet video platform like Hulu offers content that is of interest to consumers nationwide and has licensed this content for national distribution. Thus, the relevant geographic market is national.

By contrast, network neutrality proponents may classify behavior as socially harmful even if it is unlikely to monopolize the market for the affected application. In the Internet context, discrimination will often be profitable even if it does not monopolize the market for the application in question.[191] While the resulting harm may be irrelevant for antitrust law, network neutrality proposals are driven by concerns about a broader range of harms than the specific type of "harm to competition" that antitrust law is concerned with:[192] For example, exclusion allows the network provider, not the users, to choose which applications will be successful on its network. This not only distorts competition among applications on the network provider's network, but also removes an important part of the mechanism that creates innovation under uncertainty, reducing the quality of application innovation.[193] The threat of future discrimination will often reduce the incentives existing and future application providers have to innovate (not just those of the application provider that is being discriminated against) and will make it more difficult for them to get funding.[194] The resulting decline in the amount and quality of application

[189] van Schewick (2010d), pp. 231-232.

[190] Yoo (2005), pp. 71-73; Sidak (2006), p. 470, 472-473; Singer & Sidak (2007), pp. 391-392 ("[A] *local* cable modem provider with a miniscule share of *national* broadband customers lacks the ability to induce an Internet content provider from exiting the industry or even operating at an inefficient scale." Ibd., p. 391). Hemphill (2008), pp. 156-157; Hahn, Litan & Singer (2010), p. 376; Litan & Singer (2007), p. 556.

[191] van Schewick (2010d), pp. 251-255, 264-270; Frischmann & van Schewick (2007), pp. 412-416.

[192] See footnotes 30-38 above and accompanying text.

[193] See *Box 3: The Importance of User Choice* above and footnotes 53 to 54 above and accompanying text.

[194] See footnotes 229 to 230 below and accompanying text.

innovation limits the Internet's value for users and its ability to contribute to economic growth.[195] Discrimination not only deprives all Internet users of the value of future applications that would have been developed but for the threat of discrimination. It also harms the network provider's Internet access customers who cannot use the application that is being discriminated against. For applications through which users interact with others (e.g., Internet telephony or online gaming), the exclusion also harms other network providers' Internet access customers by preventing them from using the application to interact with users whose Internet access provider is blocking the application. Finally, exclusion may impair the Internet's ability to improve democratic discourse, to facilitate political organization and action, or to provide a decentralized environment for social and cultural interaction in which anyone can participate.[196] All of these harms arise even if the behavior is unlikely to monopolize the market for the applications in question.

Third, US antitrust law usually has very stringent requirements about the degree of market power in the primary market that is required for vertical exclusionary conduct to be considered problematic.[197] By contrast, network neutrality proponents are also concerned about a network provider's discriminatory behavior if that network provider does not have a dominant position in the local or nationwide market for Internet services.[198]

Fourth, under an antitrust framework, discriminatory conduct that is justified by a legitimate business purpose would be classified as socially beneficial.[199] While those who propose using an antitrust framework to distinguish between socially beneficial and socially harmful discrimination do not explain this criterion in detail, they seem to agree that conduct that is designed to increase the network provider's private efficiency should not be considered socially harmful.[200,201] For example, most proponents of an antitrust framework seem to assume

[195] van Schewick (2010d), pp. 356-361.

[196] See, e.g., Balkin (2009);van Schewick (2010d), pp. 364-365.

[197] For sources stressing the importance of market power in the market for Internet services as a prerequisite for regulatory intervention, see, e.g., Litan & Singer (2007), pp. 552-554; Yoo (2007), pp. 504, 506, 511-515; Becker, Carlton & Sider (2010). p. 505; Hahn, Litan & Singer (2010), p. 371; Sidak & Teece (2010), pp. 564-565; Hazlett & Wright (2011), Section VI., pp. 43, 44 ("[M]arket power [in the market for Internet services] is a necessary condition for such [anticompetitive] foreclosure.", ibd., p. 44); Cave, et al. (2009), pp. 1-2; AT&T Inc. (2007), pp. 66-67; Cave & Crocioni (2011), p.65. A few proponents of an antitrust framework for network neutrality do not require proof of monopoly power or of a dominant position in primary market. See, e.g., Hahn, Litan & Singer (2010), p. 367.

[198] See, e.g., van Schewick (2010d), pp. 255-264; Wu (2006c), pp. 25, 27-28; Frischmann (2012), pp. 330-332; Hogendorn (2012). See also the discussion in Section "Ban Discrimination That Is Not Disclosed" below. In the Open Internet proceeding, this position was supported by, for example, Center for Democracy & Technology (2010), pp. 9-10; Free Press (2010), pp. 45-49; Open Internet Coalition (2010), pp. 70-76; Center for Media Justice, et al. (2010), p. 23-24.

[199] See Areeda & Hovenkamp (2008), pp.183-192, para 658f (discussing the question in the context of Section 2 of the Sherman Act challenges).

[200] See, e.g., Weiser (2003), pp. 75-76; Yoo (2006); Hahn, Litan & Singer (2010), pp. 375, 376, 378.

[201] Proponents of using an antitrust framework do not provide a lot of detail about how this part of the framework would work in practice. For example, they usually do not discuss the burden of proof, or which standard should be used to decide whether the offered business justification is indeed "legitimate." The approach proposed by Yoo would assume that discriminatory behavior has a pro-competitive explanation, unless harm to consumers is proven (Yoo (2005), pp. 66-67; Yoo (2008), pp. 257-261). By contrast, Weiser assumes that discriminatory conduct is

that any discriminatory conduct that is adopted to manage congestion is pro-competitive and should be considered socially beneficial discrimination.[202] Price discrimination that is designed to recover fixed costs of network infrastructure or network innovation is often mentioned as another example of a business justification that may legitimize discriminatory conduct.[203] For those who would evaluate discriminatory conduct by network providers under an antitrust framework, the existence of an efficiency rationale ends the inquiry. The efficiencies created by the conduct do not need to outweigh any harm to competition. Nor does it matter whether there is a less restrictive alternative that may reach the same goal with less harm to competition.[204]

By contrast, network neutrality proponents often classify discriminatory behavior as socially harmful even if it is motivated by the network provider's desire to increase its own efficiency.[205] Thus, the existence of a private efficiency rationale does not automatically legitimize the behavior.

Network neutrality proponents evaluate discriminatory conduct based on its social costs and benefits. Network providers make decisions based on the conduct's private costs and benefits. As I have explained elsewhere, these decisions often diverge.[206] From the perspective of network neutrality proponents, this divergence between public interests and the network providers' private interests is a key justification for regulatory intervention. According to them, network neutrality regulation is needed precisely because what is privately efficient for network providers is not necessarily socially efficient. Under these circumstances, the fact that certain behavior is privately efficient for the network provider cannot automatically excuse the behavior.[207]

The social costs of discriminatory conduct are created by the conduct as such; they do not change depending on the network provider's motivation. If an application is being blocked, it

---

anticompetitive, unless legitimate business justification is "explained" or "offered," although it is not clear what exactly would have to be proven (see, e.g., Weiser (2003), pp. 75-76; Atkinson & Weiser (2006), pp. 57-58 (discussing discriminatory provision of Quality of Service to content providers for a fee); Weiser (2008), pp. 313-318 (same)). In cases under Section 2 of the Sherman Act, courts differ in how they assign the burden of proof that the business justification is not invalid or pretextual. See Areeda & Hovenkamp (2008), para 658f1, pp. 183-185.

[202] Yoo (2006); Brito, et al. (2010), pp. 22-23.

[203] Shelanski (2007), pp. 23-24; Weiser (2008), pp. 313, 317-318.

[204] See, e.g., Hahn, Litan & Singer (2010), pp. 375, 376, 378 in the network neutrality context. See also Areeda & Hovenkamp (2008), para 658f, pp. 189-192 (arguing against balancing of social benefits and competitive harms and against searching for a less restrictive alternative in Section 2 cases). But see United States v. Microsoft Corp., Microsoft 253 F.3d 34 (D.C. Cir. 2001), at 59 (requiring the plaintiff to "demonstrate that the anticompetitive harm of the conduct outweighs the competitive benefit" under Section 2).

[205] See, e.g., Wu (2003), pp. 168-169 (discussing price discrimination); Frischmann & van Schewick (2007) (discussing discriminatory network management); van Schewick (2010d), pp. 273-278 (discussing discriminatory pricing strategies); Frischmann (2012), pp. 348-353 (discussing discriminatory price discrimination). See also Farrell (2006) (discussing collateral damage from price discrimination that requires the exclusion of applications).

[206] van Schewick (2010d), chapter 9.

[207] By contrast, in the context of Section 2 Sherman Act, a legitimate business justification only needs to be privately efficient. See Areeda & Hovenkamp (2008), para 285f1, pp. 185-186 ("[W]hen courts speak of the business justification defense as requiring some showing of "efficiency", that term should be understood to refer to the costs or output of the monopolist itself (productive efficiency), not the market as a whole (allocative efficiency)" ibid, p. 186).

cannot reach its customers. Users will be unable to use it and the application developer and his investors will be unable to reap its benefits, whether the network provider is blocking the application to manage congestion or to exclude a competitor. Thus, the social harm – the reduction in application developers' incentives to innovate and in investors' willingness to invest, users' inability to use the Internet in the way that is most valuable to them, or their inability to participate in social, cultural or democratic discourse related to blocked content – is caused by the blocking as such, not by the motivations that are driving it.

Finally, the possibility that discriminatory behavior may increase efficiency by, for example, reducing costs or increasing performance has already been factored into the fundamental trade-off underlying calls for network neutrality regulation.[208] From the perspective of network neutrality proponents, the loss of certain short-term efficiencies from discriminatory behavior is a social cost of network neutrality rules. It is, however, the price for a system that can evolve and will remain open to new applications in the future. In other words, network neutrality rules are based on the assessment that the social benefits associated with network neutrality rules are more important than the social costs, including the loss of short-term efficiencies. Since short-term efficiency gains have already been considered and rejected as a justification for discriminatory behavior on a general basis in the fundamental trade-off underlying network neutrality regulation, the fact that a certain discriminatory conduct increases a network provider's efficiency cannot automatically justify individual instances of discriminatory behavior when they occur. After all, if legislators or regulators had deemed the loss of short-time efficiencies more important than the social benefits associated with an open, non-discriminatory Internet, they would not have adopted network neutrality rules in the first place.

All of this does not mean that network neutrality proponents will never allow discriminatory conduct that is motivated by private efficiency considerations. For example, there are circumstances under which discriminatory network management may be justified. For network neutrality proponents, however, the insight that the discriminatory conduct is designed to address a network management problem is only the beginning, not the end, of the inquiry.[209] As a result, discriminatory conduct may be considered socially harmful by network neutrality proponents, even if it were justified by a legitimate business justification and, therefore, allowed under an antitrust framework.[210]

---

[208] On this trade-off, see footnotes 40 to 45 above and accompanying text.

[209] To prevent that such an exception does not create a loophole, any exception for reasonable network management needs to be carefully defined. See Ammori (2010b) (discussing the different options for introducing loopholes into network neutrality rules) and Balkin, et al. (2009) (voicing concerns regarding ambiguities in the definition of the reasonable-network-management exception in the Open Internet NPRM). For the reasonable network management exception proposed by this paper, see *Box 21: The Exception for Reasonable Network Management* below and footnotes 516 to 528 below and accompanying text.

[210] For example, network management practices that single out specific applications or classes of applications, even though the network management problem could have been solved in application-agnostic ways, would probably be legal under an antitrust framework, but banned under the non-discrimination rule and exception for reasonable network management proposed by this paper. See Section "Ban Application-Specific Discrimination, Allow Application-Agnostic Discrimination" below.

In sum, a non-discrimination rule based on an antitrust framework will not prohibit all instances of discrimination that network neutrality proponents are concerned about. This is not an accident. Proponents of an antitrust framework view the narrow scope of antitrust-based non-discrimination rules as a "feature, not a bug."[211] For network neutrality proponents, the under-inclusiveness of an antitrust framework is a serious problem. Underlying this disagreement is a broader question about the appropriate normative framework underlying network neutrality regulation: Proponents of an antitrust framework argue that network neutrality rules should be based on the same normative framework as antitrust law. Network neutrality proponents contend that a broader normative framework is justified. This paper assumes that the decision to use the broader normative framework has been made, and evaluates alternative proposals for non-discrimination rules against this framework.

Ban Discrimination That Is Anticompetitive or Harms Users

Other proposals would ban discrimination that is "anticompetitive" or "harms users." In addition, the proposed non-discrimination rule may define certain behaviors as presumptively allowed or not allowed. For example, user-controlled prioritization may be presumptively legal; application provider-paid prioritization may be presumptively illegal. Whether a specific discriminatory behavior is anticompetitive or harms users and whether the presumptions should apply would be decided by the regulatory agency in case-by-case adjudications.

The proposal for a legislative framework on network neutrality put forward by Google and Verizon in August 2010 constitutes an example of such a rule. It prohibited "undue discrimination […] that causes meaningful harm to competition or to users," and included the rebuttable presumption that "[p]rioritization of Internet traffic would be presumed inconsistent with the non-discrimination standard."[212] (It included, however, an exception for reasonable network management that allowed network providers "to prioritize general classes or types of Internet traffic, based on latency."[213]) The FCC-led industry negotiations in the summer of 2010 seem to have focused on a non-discrimination rule of this type as well.[214]

These proposals are less specific and more ambiguous than proposals based on an antitrust framework. They use criteria that are open to interpretation without indicating which theories of harm should drive the interpretation. Instead, this decision would be made by the agency in the context of a specific adjudication. Compared to an antitrust framework, which would immediately rule out many of the cases that network neutrality proponents are concerned

[211] Hazlett & Wright (2011), p. 39 ("The FCC and net neutrality proponents often argue that that the fact that antitrust analysis might not prohibit all use of vertical contracts is a bug rather than a feature of that regime. However, that antitrust is not a 'slam dunk' can be a feature as well as a bug.")

[212] Google & Verizon (2010), p.1: "Non-Discrimination Requirement: In providing broadband Internet access service, a provider would be prohibited from engaging in undue discrimination against any lawful Internet content, application, or service in a manner that causes meaningful harm to competition or to users. Prioritization of Internet traffic would be presumed inconsistent with the non-discrimination standard, but the presumption could be rebutted."

[213] Google & Verizon (2010), p. 1.

[214] This is based on conversations with various participants in the negotiations. On the FCC-led negotiations, see, e.g., Shields (2010b); Kang (2010a); Kang (2010b).

about, these proposals could capture more of these cases under some but not all possible interpretations of the rule.

For proponents of a narrow scope of network neutrality rules, terms like "anticompetitive" or "harm to competition" are meant to evoke the standards used in antitrust analysis, where behavior is only anticompetitive if it harms competition, not just a competitor.[215] As explained above, antitrust standards would prohibit only a subset of cases that network neutrality proponents would classify as socially harmful. Under this narrow interpretation, exact outcomes would vary depending on whether the terms "anticompetitive" or "harm to competition" were used to import the full antitrust framework outlined above or only parts of that framework.

By contrast, network neutrality proponents use terms like "anticompetitive" or "harm to competition" in a looser sense that is not tied to antitrust law. To them, any discriminatory behavior that singles out specific applications or classes of applications for differential treatment distorts competition among applications or classes of applications. This harms the competitive process and, thereby, competition, by making it impossible for all applications to compete on a level playing field, without interference from network providers. It is unclear how far such an interpretation would go, but it would capture more, if not all of the cases that network neutrality proponents are concerned about than an interpretation based on antitrust law.

From the perspective of network neutrality proponents, the term "harm to users" resonates with the notion that network neutrality is designed to safeguard users' ability to use the applications of their choice and access and distribute the content of their choice without interference from network providers. There is, however, considerable uncertainty regarding the interpretation of this term. Depending on how the term is interpreted, it could capture fewer instances of discrimination than network neutrality proponents would find justified.

Consider the example of Comcast's blocking of BitTorrent. Network neutrality proponents usually agree that singling out specific applications to manage bandwidth on a network is not an acceptable form of discrimination or "reasonable network management" as long as other, application-agnostic ways of managing the network are available.[216]

An application of the rule to this case immediately raises a number of questions:

Who is a user? Singling out a specific application to manage bandwidth on a network harms the network provider's Internet access customers who want to use the application as well as the provider of the application. It is unclear, however, whether the term "harm to users" only refers to end users, or also to application and content providers.

[215] For one example, see Farber, et al. (2007).

[216] See, e.g., van Schewick (2008a), pp. 4-8; van Schewick (2010j), p. 2; Open Internet Coalition (2010), pp. 49-50 (proposing test that require the network management practice to result in "as little discrimination or preference as reasonable possible"); Center for Media Justice, et al. (2010), p. 35, 40 (same). The Canadian Radio-Television and Telecommunications Commission adopted a similar test to evaluate the Internet traffic management practices of Canadian Internet Service Providers (Canadian Radio-Television and Telecommunications Commission (2009f), asking, among other questions, whether a discriminatory network management practice results "in discrimination or preference as little as reasonably possible" ibid., para 43).

How do regulators determine whether users are harmed? Do they focus on the individual user who cannot use the Internet as she would like, or do they focus on users as a group, similar to the way antitrust law defines harm to consumers when evaluating whether a certain conduct is anticompetitive? For example, slowing down peer-to-peer file-sharing, a network provider may argue, may harm the file-sharing users and the provider of the file-sharing software, but, according to the network provider, is only done to protect the Internet experience of all the other non-file-sharing users.[217]

Does it matter that there are alternative, non-discriminatory ways of managing the network that are not similarly harmful to the users of the file-sharing software and the providers of the software, while maintaining the quality of the Internet experience for the non-file-sharing users? Network neutrality proponents usually allow discriminatory network management only if the problem cannot be solved in a non-discriminatory way,[218] but it is unclear whether a regulatory agency would read this requirement into the term "harm to users."

Finally, individual filmmakers often use peer-to-peer file-sharing applications to inexpensively distribute their creative works, as we know from the Canadian proceeding that reviewed the Internet traffic management practices of Internet service providers.[219] Non-profits often use peer-to-peer file-sharing to distribute their video contributions to political debates. Thus, peer-to-peer file-sharing applications help foster a more decentralized environment for democratic discourse and cultural production in which anybody can participate.[220] Network neutrality proponents factor the loss of these societal benefits into their evaluation of discriminatory behavior, but it is unclear whether the term "harm to users" would permit this type of consideration.

In sum, while seemingly more specific, the rule's substantive criteria are open to interpretation and do not necessarily capture the behavior that network neutrality proponents are concerned about.

Ban Discrimination That Is Unreasonable

A final set of approaches does not specify the criteria to be used in separating socially beneficial from socially harmful discrimination beyond very general terms. For example, the draft Open Internet Rules circulated by the FCC Chairman in early December 2010 banned "unreasonable" discrimination by wireline broadband Internet access providers without specifying how these

[217] See, e.g., Bell Aliant Regional Communications (2009), p.41-42, paras 85-88 (arguing that throttling peer-to-peer file-sharing applications between 4:30pm and 2am "is aimed at delivering a more positive and better experience on the network for all users" ibid., para 87); Cohen (2008), p.15 (arguing that interfering with peer-to-peer file-sharing applications "make[s] the aggregate online service better for all users and all services.")

[218] See, e.g., footnotes 440 to 443, 516 to 530 below and accompanying text.

[219] Canadian Film and Television Production Association (2009), p.12, paras 53-56 (citing concrete examples); Documentary Organization of Canada (2009), pp. 1-3 (citing concrete examples); Independent Film & Television Alliance (2009), para 8.

[220] Balkin, et al. (2009), pp. 31-38.

terms should be interpreted.[221,222] The Chairman's proposal was based on a compromise bill that had been negotiated by the Chairman of the House Committee on Energy and Commerce, Rep. Henry A. Waxman, and the Chairman of the House Subcommittee on Communications, Technology and the Internet, Rep. Rick Boucher, with the large phone and cable network providers, Internet companies, consumer groups and open Internet groups in the Fall of 2010.[223] The bill would have banned network providers from "unjustly or unreasonably discriminat[ing] in transmitting lawful traffic over a consumer's wireline broadband Internet access service."[224]

This type of rule leaves all substantive decisions about the legality of discrimination to decisions by the regulatory agency in future case-by-case adjudications, providing future decision-makers with maximum flexibility. Contrary to non-discrimination rules based on an antitrust framework, this type of proposal does not immediately rule out cases that network neutrality proponents are concerned about and makes it at least possible, but not certain that a complaint targeting behavior that network neutrality proponents deem socially harmful will be successful.

Problems with Case-by-Case Adjudication

The proposals in this group leave the substantive decision over the legality of specific discriminatory behavior to future case-by-case adjudications. The most general proposals ban

[221] The draft rules were not released publicly, but described by the Chairman in public remarks when he circulated the draft rules: "And so the proposed framework includes a bar on unreasonable discrimination in transmitting lawful network traffic." Genachowski (2010).

[222] While this language mirrors the language of the non-discrimination rule for common carriers in Section 202 of the Communications Act (47 U.S.C. 202 (a), full text below), it is not clear whether and, if so, how this would affect the interpretation of the rule. Motivated at least in part by heavy resistance from network providers, the FCC made a deliberate decision not to reclassify Internet access services as telecommunications services, which would have made Section 202 immediately applicable. The decision not to reclassify could be interpreted as a deliberate decision against the substantive framework provided by Title II. Even if interpreters were willing to look to Title II's common-carrier non-discrimination rule for guidance on how to interpret terms in a network-neutrality non-discrimination rule enacted under Title I of the Communications Act, it is unclear whether precedents developed for telephony services would be applicable to Internet access services. Beyond this general objection, lawyers would find ample opportunity to fight over the relevance of specific precedents. See, for example, the exchange between Harold Feld, legal director of the public interest group Public Knowledge, and lawyers for AT&T regarding the consequences of reclassifying Internet access services as a telecommunications service for the treatment of paid prioritization under the resulting Title II regime. (Feld (2010c); Feld (2010a); Quinn (2010); Hultquist (2010)). See also Lyons (2013).

47 U.S.C. 202: (a) "It shall be unlawful for any common carrier to make any unjust or unreasonable discrimination in charges, practices, classifications, regulations, facilities, or services for or in connection with like communication service, directly or indirectly, by any means or device, or to make or give any undue or unreasonable preference or advantage to any particular person, class of persons, or locality, or to subject any particular person, class of persons, or locality to any undue or unreasonable prejudice or disadvantage."

[223] Jerome (2010b); Waxman (2010). Not all participants in the negotiations supported the final proposal (Jerome (2010a); Eggerton (2010)). It was dropped when the Republican members of the Energy and Commerce Committee refused to support the bill (Hart (2010a); Waxman (2010)). Chairman Waxman later filed the draft bill in the Open Internet docket (Draft Bill To Amend Title I of the Communications Act of 1934 to Provide for Internet Openness, and for Other Purposes (2010)). On the Waxman proposal's influence on the Chairman's proposal, see footnote 542 below and footnotes 554 to 555 below and accompanying text.

[224] Draft Bill To Amend Title I of the Communications Act of 1934 to Provide for Internet Openness, and for Other Purposes (2010), Section 12(a)(1)(B).

"unreasonable discrimination," but do not provide any guidance on how to distinguish socially beneficial from socially harmful discrimination, leaving both the development of substantive criteria as well as their application to the specific behavior under consideration to future decision-makers. While proposals that prohibit discrimination that "creates meaningful harm to competition or to users" seem more specific, they are afflicted with the same problem. The outcome of any adjudication depends entirely on how these ambiguous terms would be interpreted, with different interpretations leading to radically different outcomes. Other non-discrimination rules evaluate discriminatory conduct after the fact using multiple factors without specifying how the factors relate to each other. Here, the outcome of specific adjudications depends not only on how future decision makers interpret and apply those factors, but also how they weigh the different factors against each other. The non-discrimination rule proposed by the FCC in its May 2015 Notice of Proposed Rulemaking is an example of such a rule.[225]

These kinds of case-by-case approaches create considerable social costs:[226]

First, case-by-case approaches fail to provide much-needed certainty for industry participants. While individual adjudications may reduce the amount of uncertainty over time, it is unclear whether and how fast useful precedents will emerge. (See *Box 12: Will Individual Adjudications Reduce Uncertainty Over Time?* below.)

Under the proposals discussed above, network providers do not know which forms of network management are acceptable. For example, it is unclear whether, and if so, which forms of Quality of Service would be considered socially beneficial in future applications of the rule. It seems rather unlikely that network providers would make the investment needed to introduce Quality of Service in their Internet access networks if that investment could subsequently be made moot if a regulator, following a complaint, declared the practice socially harmful.[227] By contrast, the more nuanced rules described below would clearly allow certain, though not all forms of Quality of Service. Thus, under a case-by-case approach, network providers may refrain from deploying network technology that would have been clearly legal under one of the more nuanced rules discussed below. The resulting lack of evolution of the network infrastructure harms innovation in applications that need Quality of Service and deprives users of the benefits associated with the emergence of these applications.

[225] The non-discrimination rule for fixed broadband Internet access proposed by the FCC would ban "commercially unreasonable discrimination." Whether certain discriminatory conduct is commercially unreasonable would be determined after the fact on a case-by-case basis, using a number of factors that have yet to be specified, taking into account the totality of circumstances. Federal Communications Commission (2014), pp. 39-50, paras 110-141, pp. 58-61, paras 161-169.

[226] The relative costs and benefits of laws that specify legal commands before individuals act (i.e., when the law is promulgated) ("rules") and laws where legal commands are specified after individuals have acted (i.e., in the context of adjudication) ("standards") as well as the costs and benefits of proceeding by adjudication rather than rulemaking are discussed by three bodies of literature: the literature on rules vs. standards, the literature on rulemaking vs. adjudication in administrative law, and the literature on per se rules vs. rule of reason in antitrust law.

[227] National Cable & Telecommunications Association (2007), p. 15; Verizon & Verizon Wireless (2007), pp. 44-45.

More generally, some research and anecdotal evidence suggest that in the broadband context, certainty regarding the regulatory framework and its stability over time may be more important for network investment than the substance of the regulatory decision.[228]

In a network that can identify applications and control their execution, application developers who must decide whether to realize their innovative ideas and investors who consider funding them face the fundamental risk that the network may discriminate against the application at any time, which would reduce the affected application provider's ability to reap the benefits associated with her innovation. Thus, the threat of discrimination reduces application developers' incentives to innovate and their ability to get funding.[229,230] Network neutrality rules aim at mitigating that problem by providing application developers and their investors with certainty that they will not be discriminated against. A case-by-case approach falls short of this goal. Innovators and their investors will not know in advance if and against which network provider conduct they are protected, because this decision will only be made after discriminatory conduct has occurred. If the application is discriminated against, its chances with users are harmed immediately, and this harm persists while the application provider goes through a long and costly process to reach a regulatory decision on the discriminatory behavior in question. In markets where first-mover advantages are important, the temporary disadvantage may be sufficient to tip the competition against the affected application. Moreover, venture capitalists and other investors fund start-ups so that these companies can build their products and better meet the needs of their users. Paying lawyers and economists to clarify how to interpret an ambiguous non-discrimination rule in order to allow the application to reach its customers is not how investors would like their funds to be used. Thus, this type of non-discrimination rule does not sufficiently protect users and application developers against actual discrimination and fails to remove the threat of discrimination as a factor that affects application developers' and innovators' decisions about innovation.[231]

[228] See, e.g., Wernick (2007), pp. 129-130, 158-185; Verizon & Verizon Wireless (2007), pp. 44-45.

[229] See, e.g., Lessig (2008b), pp. 7-8; van Schewick (2010d), pp. 270-273.

[230] As I know from many conversations with entrepreneurs and investors, the threat of discrimination reduces entrepreneurs' ability to get funding today. For two publicly documented examples, see van Schewick (2008b), p. 2 and the letter from the founders of the online video company Zediva to the FCC (Srinivasan & Gupta (2010), pp. 1-2). See also Lessig (2008b), pp. 7-8.

[231] Of course, even a rule-based approach that specifies in advance which differential treatment is and is not acceptable cannot provide absolute certainty that discriminatory behavior that violates the rule will never occur. In the presence of clear rules, however, network providers will know what behavior is not allowed and may prefer complying over risking detection, enforcement and fines (the network provider's exact calculus depends on the likelihood of detection, the agency's willingness to enforce the rule and on the sanctions associated with a rule violation). In addition, since enforcing a clear rule is easier, faster and less costly than engaging in the type of case-by-case adjudication described above, the overall harm to an application developer and to users if discrimination occurs is smaller than under a case-by-case regime.

Box 12

Will Individual Adjudications Reduce Uncertainty Over Time?

Over time, individual adjudications may clarify the interpretation of the standard and its application to specific behavior, reducing uncertainty.[232] Whether future adjudications manage to reduce uncertainty in a meaningful way depends on a variety of factors: First, network providers need to be willing to engage in discriminatory conduct and take the risk of being faced with a complaint and having the behavior declared socially harmful. If network providers do not engage in a particular practice (e.g., if they do not deploy Quality of Service in their networks), there is no basis for a complaint and the legality of the practice will never be determined. Second, contrary to a rule that clearly specifies which behavior is and is not allowed, an adjudicatory regime puts the burden on a particular party to bring a complaint that will allow the uncertainty to be resolved. Third, future adjudicators may not be any more willing than the current legislator or regulator to do more than absolutely necessary to resolve the case under consideration. Narrow decisions that are deliberately tied to the facts of the specific case and refuse to elaborate broader principles may not provide meaningful guidance for future cases.[233] Thus, it is unclear whether and how quickly useful precedents will emerge. In the meantime, the costs associated with the uncertainty persist.[234] Moreover, as set out in more detail below, the substantive principles emerging from case-by-case adjudications are less likely to adequately protect the values and actors that network neutrality rules are designed to protect.

Second, case-by-case approaches create high costs of regulation.[235] Each adjudication requires detailed investigations into the facts of the case and invites protracted and resource-intensive fights over the interpretation of the rule. Precedents established through adjudication may not necessarily be binding on other industry actors.[236] Their applicability may also be limited by the facts of the case.[237] As a result, subsequent cases may need to be fully adjudicated, even if they are based on similar facts, with network providers arguing that the facts of their case differ from the precedent in relevant ways. For example, when the FCC ordered Comcast to stop interfering with BitTorrent and adopt application-agnostic ways of managing congestion,[238] the Commission based its decision on three different rationales: First, the specific practice used by Comcast – sending RST packets to terminate BitTorrent connections – was quite questionable and violated the Internet Engineering Task Force's standards for the operation of the Transport Control

[232] See, e.g., Yoo (2009), p. 82.

[233] See, e.g., Kaplow (1992), pp. 612-613 (discussing courts' tendency to issue narrow decisions that minimize or postpone the creation of precedents); Sunstein (1995a) (discussing the dynamics that give rise to this phenomenon).

[234] See, e.g., Yoo (2009), p. 82 (noting the uncertainty surrounding the FCC's interpretation of "reasonable network management" until enough precedents exist). See generally Kaplow (1992), pp. 612-614 and Kaplow (1999), p. 512 (noting the costs associated with waiting for precedents to resolve an open legal question instead of resolving the question through an earlier announcement of a rule).

[235] See, e.g., Kaplow (1999), p. 510; Pierce (2010), pp. 497-508.

[236] Magill (2004), p. 1394, 1396; Pierce (2010), pp. 498-499.

[237] Magill (2004), p. 1396; Pierce (2010), p.499, 500.

[238] Federal Communications Commission (2008), pp. 32-33, paras 54-55 and p. 34, para 59.

Protocol.[239] Second, the discriminatory practice, which singled out BitTorrent and other peer-to-peer file-sharing applications for differential treatment, was not narrowly tailored to Comcast's stated goal of managing congestion.[240] Third, Comcast had not disclosed the use of the practice to its Internet access customers.[241] The Order did not explain whether each of these factors alone would have made the network management "unreasonable", or whether the Commission's decision was based on the confluence of these factors, providing ample room for network providers to distinguish their case on the basis that their behavior violated only one, but not all the criteria used in the Comcast case.[242]

Finally, in the context of network neutrality, case-by-case approaches are less likely than rule-based approaches to adequately protect the values and actors that network neutrality rules aim to protect.

Case-by-case approaches provide an advantage to well-financed actors and tilt the playing field against those – end users, low-cost application developers and start-ups – who do not have the resources necessary to engage in extended fights over the legality of specific discriminations in the future.[243] Network providers and large application providers can conduct fact-intensive investigations, pay lawyers, economists and other experts to engage in the fight over the correct interpretation and the application of the rule at the regulatory agency and, later, in the courts, and employ lobbyists to organize support for their position in Congress or at the White House. End users, low-cost application developers and start-ups lack these resources. Thus, adjudications will likely be systematically biased against their interests. They are, however, some of the key groups that network neutrality rules are intended to protect.[244]

Decisions in individual adjudications will often be driven by the specific facts of the case. A sympathetic party or a limited fact-pattern that does not illuminate all relevant aspects of the underlying problem may distort the decision-maker's view of the underlying policy issues in a way that a more general analysis of the issues in the context of a rulemaking proceeding may not.[245] For example, as in the FCC's investigation of Comcast's blocking of BitTorrent, debates over the reasonableness of network management practices arose first in the context of discriminatory treatment of peer-to-peer file-sharing applications. Most people have heard of BitTorrent and other peer-to-peer file-sharing applications as tools for illegal file-sharing. They

[239] Federal Communications Commission (2008), pp. 27-28, paras 45-46.

[240] Federal Communications Commission (2008), pp. 28-31, paras, 47-51.

[241] Federal Communications Commission (2008), pp. 31-32, paras 52-53.

[242] The Comcast Order was vacated by the United States Court of Appeals for the District of Columbia Circuit in April 2010. Comcast Corp. v. FCC, 600 F.3d 642 (D.C. Cir. 2010).

[243] See generally, e.g., Croley (1998), pp. 120-124, 128; Sunstein (1995b), p. 977("[C]ase-by-case judgments systematically favor the well-to-do."). In the network neutrality context, see van Schewick (2010e), p. 5; Burnham (2010b).

[244] On the importance of low cost innovators, see van Schewick (2010d), pp. 204-213, 300-308, 334-345. On new entrants and start-ups, see van Schewick (2010d), pp. 319-334. For a short version of the argument, see van Schewick (2010h), pp. 3-5.

[245] Magill (2004), p. 1396; Pierce (2010), pp. 496-497. See also Kaplow (1992), p. 609 (discussing this phenomenon with respect to the decisions of courts compared to legislative decisions).

do not know that peer-to-peer file-sharing applications have many legal and socially valuable uses. For example, at the time of Comcast's blocking of BitTorrent, established content providers such as the BBC, Showtime, the History Channel, MTV Media Networks, 20th Century Fox, Paramount and many others were distributing their video content online through services that utilized the BitTorrent protocol.[246] Developers of open source applications such as the Linux operating system or OpenOffice or game providers such as Blizzard Entertainment, the company behind World of Warcraft, employ peer-to-peer file-sharing applications to distribute their software or software updates.[247] Peer-to-peer file-sharing applications foster a more decentralized environment for the creation and distribution of creative works by allowing independent filmmakers to sidestep traditional, more centralized distribution channels and distribute their films directly to the public.[248] Internet video applications based on peer-to-peer protocols like the Miro video player let a diverse set of actors distribute their videos on a wide range of subjects, providing an important outlet for free speech.[249] Still, based on the inaccurate perception that applications like BitTorrent are primarily used for illegal file-sharing, regulators, members of Congress or the White House may be more reluctant to side with complaints against network management practices that single out these applications. After all, who wants to side with "pirates"? More generally, the question at the core of the debate over reasonable congestion management – who should prioritize among competing uses at times when people most want to use the network – may receive more attention and a more balanced assessment in a general rulemaking than in an adjudication involving peer-to-peer file-sharing applications. Adjudications focused solely on peer-to-peer file-sharing applications foster the general perception that network providers engage in congestion management to protect socially valuable applications from the bandwidth demands of applications that have little social value, providing little reason to question network providers' role as benevolent steward of the platform. By contrast, a more general analysis of network management practices would broaden the focus to include attempts to limit the use of other applications, for example of streaming video applications, during times of congestion. In 2009, for example, British Telecom (BT) restricted the bandwidth available to the BBC iPlayer and other streaming video applications to 986 kilobytes/sec in a particular version of BT's broadband service.[250] Many people like to use streaming video applications like Hulu or Netflix in the evening, when the network is most congested. In North America, Netflix traffic now makes up 34% of downstream traffic on fixed broadband networks during peak times.[251] As a result, in a generalized rulemaking that also considers limits on applications other than peer-to-peer file-sharing applications, the sympathy of decision-makers and observers will be more evenly distributed among restricted and unrestricted uses of the network. At the same time, streaming applications, which compete with

[246] Free Press & Public Knowledge (2007), pp. 17-18.

[247] Linux Tracker (2012); OpenOffice (2012); Free Press & Public Knowledge (2007), pp. 18-19.

[248] Canadian Film and Television Production Association (2009), p.12, paras 53-56, 58; Documentary Organization of Canada (2009), pp. 1-3, 5.

[249] Balkin, et al. (2009), pp. 31-38.

[250] Cellan-Jones (2009).

[251] Sandvine (2014), p. 5.

network providers' traditional video offerings, bring the potential gap between network providers' and users' interests into sharp relief,[252] making the argument more convincing that users, not network providers are in the best position to decide how the network should be used, whether there is congestion or not. For all these reasons, an individual adjudication focused on network management practices singling out peer-to-peer file-sharing is more likely to result in a decision that grants network providers broad discretion in managing congestion than a general rulemaking. At the same time, the precedent set by the adjudicatory decision may make it more difficult to limit network providers' discretion when congestion management practices arise that target other uses of the network.

More generally, adjudicators who need to decide whether a certain discriminatory behavior should be allowed as part of an adjudication will be less likely to have access to the full set of relevant facts and arguments than public actors trying to distinguish socially beneficial from socially harmful discrimination as part of a rulemaking.[253] Contrary to rulemakings, adjudications are adversarial proceedings, with procedural rules that make it more difficult for other interested actors to participate. This limits the range of actors from which the adjudicator will receive input.[254] This is particularly problematic in the context of network neutrality rules, where any decision over the legality of discriminatory behavior is likely to have far-reaching implications for users, application providers, their investors and network providers who are not directly subject to the discriminatory practice under consideration. Moreover, network neutrality rules are designed to protect, among others, the interests of users as well as of current and future innovators and entrepreneurs. As large groups with diffuse interests, they face well-documented challenges in organizing and representing their interests, which makes it more difficult for them to participate and be heard in any type of legislative or regulatory proceeding.[255] Adversarial proceedings increase these challenges.[256] For example, entrepreneurs are often reluctant to speak out on network neutrality because they fear retaliation by network providers.[257] They may be even more reluctant to do so in the context of an adjudication which is directed against a specific network provider. Also, it may be easier to mobilize users and entrepreneurs once, in the context of a rulemaking, than again and again for individual adjudications. Users or entrepreneurs may not only find it difficult to understand how a specific adjudication may affect them. Like public decision-makers, they may also be subject to biases or

[252] See, e.g., Canadian Film and Television Production Association (2009), pp. 13-14, paras 60-67; U.S. Department of Justice (2011), pp. 11, 14-20 (discussing the gap and citing "[m]any internal Comcast documents" showing that Comcast views online video distributors as a competitive threat to its traditional cable video distribution offerings, ibid., p.19).

[253] See generally Magill (2004), p.1396, 1446; Pierce (2010), pp. 496-497, 501.

[254] See generally, Croley (1998), pp. 116-177, 120-124, 128, 148; Magill (2004), p. 1391, 1396; Pierce (2010), pp. 496-497, 501.

[255] Olson (1980), pp. 2, 9-16, 35; Croley (1998), pp. 126-142; Bagley & Revesz (2006), pp. 1287-1290; Barkow (2010), pp. 21-24..

[256] Croley (1998), pp. 120-124, 128 (discussing the costs of participating in rulemaking and adjudication and the impact on relative levels of participation of public interest groups in these types of agency decision-making and citing empirical studies supporting his analysis).

[257] van Schewick (2010l), Section "This is one example of many."

intuitive reactions resulting from an adjudication's specific fact patterns.[258] For example, a user who does not use BitTorrent and does not engage in illegal file-sharing may fail to grasp the importance of an adjudication focused on network management practices targeting peer-to-peer file-sharing. Entrepreneurs offering streaming video applications that do not use peer-to-peer protocols may have the same reaction. For all these reasons, users and entrepreneurs may be less willing to get involved in specific adjudications than in a general rulemaking, depriving the decision-maker of input from important stakeholders.

Finally, an ex-ante regime is better suited to the consideration of the very fundamental values at stake than case-by-case adjudications. Network neutrality rules are based on very general trade-offs among competing values.[259] Network neutrality rules foster application innovation, protect user choice, and preserve, among other things, the Internet's ability to foster democratic discourse, all of which creates social value. They limit the evolution of the network's core to some extent, limit network providers' ability to realize all potential efficiency gains or optimize the network in favor of the applications of the day, reduce network providers' profits and, like all regulation, need to be administered and enforced, all of which creates social costs. Thus, there is a trade-off that regulators need to resolve. An ex-ante rule that specifies what behavior is and is not allowed resolves this trade-off for all future cases at once, in favor of the social benefits. If the legality of discriminatory behavior is decided case-by-case instead, it is more likely that decisions will deviate from this general trade-off and allow discriminatory behavior than under a rule that makes this decision *ex ante*. This is because the adjudicator's decision will be affected by several well-known cognitive limitations and biases.[260]

While the costs of banning the practice will be immediately apparent (e.g., the network provider cannot manage its network in a certain discriminatory way), the current and future benefits associated with a ban will be less clear. While the discriminatory practice immediately harms the provider and the existing users of the affected application, the value of a specific application often only becomes apparent over time. Thus, the immediate cost of the discriminatory practice (or the immediate benefit of banning it) may be difficult to quantify. Determining the future benefits of banning the discriminatory practice is even more difficult. We do not know which applications will never be developed because innovators and investors are concerned about the threat of discrimination, so their social value cannot be determined, either.[261]

Moreover, an adjudicator is likely to underestimate other negative consequences of allowing a deviation from the general nondiscrimination rule in the particular case under

[258] See footnotes (and accompanying text) 245 to 252, 262 to 272, 294 to 295 below.

[259] For a detailed discussion of this trade-off, see van Schewick (2010d), pp. 355-371. For a short overview, see the discussion in Section "A Framework for Evaluating Network Neutrality Rules" above.

[260] van Schewick (2010d), p. 78 (making the same argument in the context of deviations from the broad version of the end-to-end arguments). See also Solum & Chung (2004), pp. 854-865 (making a similar argument in a slightly different context).

[261] For a more detailed description of the problem with pointers to the literature, see van Schewick (2010d), pp. 77-78, 374-375.

consideration. Often, it takes a while to recognize the negative consequences of a specific discriminatory practice (beyond any reduction in incentives to innovate due to the threat of discrimination). This problem may be particularly pronounced for an adjudicator who lacks technical expertise.[262] For example, network management practices that single out specific applications or classes of applications for negative treatment may motivate the designers of the affected applications to adopt techniques to evade detection:[263] Applications that are the target of discriminatory network management practices and others that want to avoid being targeted in the future often choose to encrypt their communications across the network.[264] The increase in encryption has motivated some operators to slow down all encrypted traffic, which in turn hurts legitimate traffic that is encrypted for security reasons.[265] Widespread use of encryption also complicates network analysis, planning and security.[266] Similarly, Comcast's old, discriminatory method of managing congestion – sending spoofed RST packets to terminate certain peer-to-peer file-sharing connections – used certain types of TCP packets in a non-standard way. Once such a practice emerges, programmers can no longer rely on standards to determine how their software should respond to an RST packet, which considerably complicates protocol and application design.[267] Thus, allowing only a single discriminatory network management practice (e.g., one targeting peer-to-peer file-sharing applications) may have significant unintended negative consequences. Beyond that, several small deviations may quickly add up to create big roadblocks for innovation.[268] For example, while application developers may be able to adopt their application to one network provider's idiosyncratic discriminatory network management practice, the costs of adopting their application to the network management practices of more than a few providers will quickly become prohibitive.[269] Thus, an adjudicator's focus on a single practice whose exact effects may yet be unknown is likely to lead him to underestimate both the isolated effect of the practice as well as its interactions with other current or future deviations from non-discriminatory network management. By contrast, decision-makers in a general rulemaking can take a broader view that takes account of cumulative effects and generalizes from past experiences.[270]

---

[262] van Schewick (2010d), p. 78; Solum & Chung (2004), pp. 859-860.

[263] The resulting arms-race is described in detail by Lehr, et al. (2007).

[264] Bode (2006); Blumenthal & Clark (2001), pp. 77-78, 95; Lehr, et al. (2007), pp. 627-628.

[265] For example, in 2007, Canadian cable provider Rogers reportedly throttled all encrypted traffic. Observers suggested that this negatively affected students and professors at the University of Ottawa who were Rogers' customers and were checking their e-mail from off-campus; off-campus access to e-mail is encrypted for security reasons. Geist (2007). In response to the allegations, Rogers declared that it was not degrading encrypted traffic and that its own tests had not shown any performance problems for the most common encrypted applications. Engelhart (2007).

[266] Blumenthal & Clark (2001), pp. 77-78, 86; Lehr, et al. (2007), p. 635.

[267] See, e.g., Eckersley, von Lohmann & Schoen (2007). Applications subject to discriminatory network management may also masquerade as other applications that are not subject to the discriminatory practice, which makes it more likely that these other applications will be inadvertently caught by the discriminatory network management measures.

[268] van Schewick (2010d), p. 78, 374; Solum & Chung (2004), pp. 855-856, 859-860.

[269] Lessig (2008a), p. 2;

[270] Solum & Chung (2004), pp. 855-856, 859-860.

Finally, research in behavioral economics suggests that individuals tend to systematically undervalue future benefits, discounting them more than rational discounting would suggest.[271] Uncertainty about future benefits aggravates this bias.[272] Thus, in weighing the immediate benefits of allowing the discriminatory practice against the future, uncertain benefits of a ban, an adjudicator will disproportionately discount the future benefits.

For all these reasons, deciding whether to allow discrimination on a case-by-case basis makes it more likely that discrimination will be allowed than under an ex-ante rule that resolves the above trade-off for all future cases at once.

In spite of these considerable social costs, the strategic incentives of legislators or regulators who consider adopting network neutrality rules and of the big stakeholders on both sides of the debate are aligned in favor of such a scheme.[273] Stakeholders cannot agree which discriminatory behavior is acceptable today; it is unlikely that they will be able to do so in the future. There are large, well-financed entities on both sides of the network neutrality debate. Any substantive decision would take on either the large, well-financed, well-organized and politically influential network providers (e.g., in the US, AT&T, Verizon, Comcast and Time Warner) or big Internet application, content or service providers such as Google or Amazon. Under these circumstances, adopting a very general or ambiguous non-discrimination rule today constitutes an attractive compromise, since the controversial question is not decided one way or the other.[274]

The legislator or regulator can reap any immediate benefits associated with "adopting network neutrality rules,"[275] while avoiding the immediate political costs of taking on powerful interests on one side of the debate.[276] While an ambiguous or general non-discrimination rule that is applied case-by-case is more difficult and costly to apply and enforce in the future, these

[271] This bias is known as hyperbolic discounting. See generally, e.g., Loewenstein & Prelec (1992); Frederick, Loewenstein & O'Donoghue (2002); Read (2004). See also van Schewick (2010d), p. 78, 374 (discussing how hyperbolic discounting affects decisions to deviate from the broad version of the end-to-end arguments).

[272] van Schewick (2010d), p. 78, 374. See generally, e.g., Read (2004), p.431.

[273] See generally, e.g., Eskridge (1988), pp. 288-289; Kaplow (1992), p. 609 and footnotes 141 and 143; Sunstein (1995b), p. 973, 1004-1005.

[274] See generally, e.g., Eskridge (1988), p.288-289; Sunstein (1995b), p. 1014 ([S]tandards are more likely to be the basis for decision when opposing interests have roughly equivalent power in the lawmaking body, and when they are equally willing to take their chances with a bureaucracy or judge.")

[275] For example, by adopting network neutrality rules, the FCC fulfilled a campaign promise by President Obama, who had promised to adopt network neutrality rules if he got elected (Obama for America (2007)) and satisfied the democratic leadership in Congress, which also had supported network neutrality regulation (Feld (2010b); Ammori (2010a)). The adoption of network neutrality rules in the US was also supported by large Internet companies (e.g., Amazon, eBay, Facebook or Google), entrepreneurs and investors (Kopf, et al. (2010); Akhund, et al. (2010)), public interest groups (Soule, et al. (2009)) and a large grass roots movement (Karr (2010)).

[276] See generally, e.g., Kaplow (1992), p. 609 and footnotes 141 and 143; Sunstein (1995b), p.973, 1004, 1013 ("When lawmaking is separate from law-interpretation and law-enforcement, many of the costs of producing clarity *ex ante* will be faced by lawmakers themselves, whereas many of the costs of producing clarity *ex post* will be faced by others." Ibid, p. 1004) (all discussing the cost-benefit calculus discussed in the text in the context of a legislature's choice between rules and standards).

costs will not be borne by the entity adopting the rule:[277] If the non-discrimination rule is adopted through legislation, it will most likely be enforced by a regulatory agency (e.g., in the US, by the FCC). Even if the non-discrimination rule is adopted by a regulatory agency such as the FCC through an administrative rulemaking, it may be enforced by future members of the agency (e.g., in the case of the FCC, by future Commissions) or by another entity within the agency (e.g., the FCC's Enforcement Bureau). The social costs of this type of rule will not be borne by the entity adopting the non-discrimination rule, either.[278]

Big stakeholders support this type of non-discrimination rule because each side can claim a win (or, at least, a non-loss), and gets a second chance to influence the ultimate decision over the legality of specific practices in the context of individual adjudications in the future. While adjudications are costly, big stakeholders have the resources to play the case-by-case game and prevail in future adjudications. Given these incentives, it is not surprising that the proposals for a general or ambiguous non-discrimination rule described above emerged from industry negotiations at the FCC[279] and in Congress,[280] or, as the Verizon-Google legislative framework proposal, from direct negotiations between two big stakeholders on opposite sides of the debate.[281]

So far, the discussion has focused on proposals for general or ambiguous non-discrimination standards. Compared to these proposals, case-by-case approaches based on an antitrust framework provide considerably more guidance on how to evaluate discriminatory behavior. Even under an antitrust framework, the outcome of specific cases depends on the exact interpretation of the framework and on its application to the facts. In addition, the outcome depends on many facts specific to the case (e.g., the network provider's market share in the nationwide market for Internet access services, the existence and size of economies of scale, and the cost disadvantage associated with operating at a less than efficient scale), some of which are difficult and costly to prove.[282] As a result, an antitrust framework shares many of the

[277] See generally, e.g., Sunstein (1995b), p.973, 1004, 1013 ("When lawmaking is separate from law-interpretation and law-enforcement, many of the costs of producing clarity *ex ante* will be faced by lawmakers themselves, whereas many of the costs of producing clarity *ex post* will be faced by others." Ibid, p. 1004)

[278] See generally, e.g., Sunstein (1995b), p.973.

[279] The FCC-led industry negotiations included Google, Skype, the Open Internet Coalition (which, at the time, represented, among others, Google, Skype, Free Press, Public Knowledge, Amazon and Sony Electronics), AT&T, Verizon and the National Cable and Telecommunications Association (NCTA), which represents larger cable operators. Kang (2010a).

[280] The FCC Chairman's draft Open Internet Rules were based on a proposal for a network neutrality bill that had been negotiated by Rep. Henry A. Waxman, the Chairman of the House Committee on Energy and Commerce, and Rep. Rick Boucher, the Chairman of the House Subcommittee on Communications, Technology and the Internet with the large phone and cable network providers, Internet companies, consumer groups and open Internet groups in the Fall of 2010. Not all participants in the negotiations backed the final proposal. See footnote 223 above and, on the influence of the draft bill on the Chairman's draft rules, footnote 542 below and footnotes 554 to 555 below and accompanying text.

[281] At the end of the third quarter in 2011, Verizon was the fourth largest broadband Internet access service provider, the second largest telephone broadband Internet access provider and the largest wireless carrier in the United States. Leichtman Research Group (2011).

[282] See Areeda, Hovenkamp & Solow (2007), pp. 44-48, para 408b (arguing that precise measurement of economies of scale is very difficult, if not impossible) and pp. 54-58, para 408d (arguing that "[a]ntitrust rules requiring judicial

social costs described above with the more general and ambiguous proposals. In particular, the uncertainty about the legality of specific discriminatory conduct is not resolved until after the discrimination has occurred. Since the outcome of an adjudication depends on the specific facts of the case, the same practice may be legal for some providers, but not others, or with respect to some applications, but not others. Thus, prior adjudications will not necessarily remove the uncertainty. Finally, like the general or ambiguous non-discrimination rules discussed above, a non-discrimination standard based on an antitrust framework creates high costs of regulation, tilts the playing field against those who do not have the resources to engage in lengthy and costly fights over the legality of discrimination, and usually limits the ability of interested third parties to participate in the adjudication.

While the social costs of a non-discrimination standard based on an antitrust framework are similar to general or ambiguous non-discrimination standards, it is less attractive as a compromise solution. The substantive criteria clearly exclude many cases that network neutrality proponents are concerned about, making the standard much more favorable to network providers than to network neutrality proponents.

MORE NUANCED RULES

A final group of proposals would adopt more nuanced rules that specify in advance which differential treatment is and is not allowed. Like the standards-based approaches discussed above, these proposals recognize that some forms of discrimination are socially beneficial, while others are socially harmful. Contrary to those approaches, however, they define in advance what constitutes acceptable and unacceptable discrimination to avoid the social costs associated with leaving the decision about specific discriminatory conduct to future case-by-case adjudications.

Formal Approaches: Ban Discrimination That Is Not Disclosed

The first set of approaches in this group bans discrimination that is not disclosed, distinguishing between socially beneficial and socially harmful practices using a formal criterion. Alternatively, a network neutrality regime might allow blocking or discrimination, but require Internet service providers to disclose any blocking or discrimination that occurs.[283] In January 2014, the Court of Appeals for the D.C. Circuit struck down the Open Internet Order's rules against blocking and discrimination, but upheld the disclosure rule.[284] Thus, until the FCC adopts new network

---

measurement of scale economies should be avoided" since "[c]ourts are simply not up to the task of deciding ex jure such questions as […] that a practice should be condemned because it denies a firm scale economies except in the clearest of cases" ibid., p. 54).

[283] See, e.g., Cave, et al. (2009), pp. 2-3; AT&T, et al. (2008), pp. 1, 3; Mayo, et al. (2010); Faulhaber (2009), pp. 762-766; Faulhaber & Farber (2010b), pp. 315-316.

[284] Verizon v. FCC, No. 11-1355 (D.C.Cir. 2014), p. 62, available online at http://www.cadc.uscourts.gov/internet/opinions.nsf/3AF8B4D938CDEEA685257C6000532062/$file/11-1355-1474943.pdf.

neutrality rules, the current network neutrality regime in the US constitutes an example of this approach.

In 2009, the European Union adopted this approach following the review of its regulatory framework for telecommunications services.[285] The European Universal Service Directive neither requires network providers to impose restrictions on users' use of applications, nor does it prevent them from doing so.[286] It does, however, require Internet access service providers to inform their customers about any limits on access to or the use of services and applications, about any traffic management measures and their impact on quality of service and about minimum quality of service levels. This information must be disclosed in the terms of the contract and when practices change.

This approach is based on the idea that if a network provider discriminates against an application that users would like to use, users can switch to another network provider who does not discriminate against the affected application. The threat of switching, proponents of this approach assume, will discipline providers.[287]

---

[285] See Articles 20 and 21 Directive 2002/22/EC of the European Parliament and of the Council of 7 March 2002, as amended by Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009 (Universal Service Directive).

[286] See Preamble 29 Universal Service Directive:

> "Directive 2002/22/EC (Universal Service Directive) neither mandates nor prohibits conditions imposed by providers, in accordance with national law, limiting end-users' access to and/or use of services and applications, but lays down an obligation to provide information regarding such conditions."

Instead, the Directive leaves it to the member states to implement further rules in this respect:

> "Member States wishing to implement measures regarding end-users' access to and/or use of services and applications must respect the fundamental rights of citizens, including in relation to privacy and due process, and any such measures should take full account of policy goals defined at Community level, such as furthering the development of the Community information society." (Preamble 29 Universal Service Directive).

At the same time, the Universal Service Directive as well as the Framework Directive stress that users should be able to choose how they want to use the Internet:

> "End-users should be able to decide what content they want to send and receive, and which services, applications, hardware and software they want to use for such purposes, without prejudice to the need to preserve the integrity and security of networks and services. A competitive market will provide users with a wide choice of content, applications and services. National regulatory authorities should promote users' ability to access and distribute information and to run applications and services of their choice, as provided for in Article 8 of Directive 2002/21/EC (Framework Directive)." (Preamble 28 Universal Service Directive).

Art. 8 Framework Directive (Directive 2002/21/EC of the European Parliament and of the Council of 7 March 2002, as amended by Directive 2009/140/EC of the European Parliament and of the Council of 25 November 2009) explicitly requires national regulatory authorities to promote this goal:

> "2. The national regulatory authorities shall promote competition in the provision of electronic communications networks, electronic communications services and associated facilities and services by inter alia: […] (b) ensuring that there is no distortion or restriction of competition in the electronic communications sector, including the transmission of content. […] 4. The national regulatory authorities shall promote the interests of the citizens of the European Union by inter alia: […] promoting the ability of end-users to access and distribute information or run applications and services of their choice."

[287] See, e.g., European Commission (2007), p. 91; European Commission (2011), p. 4.

Disclosure can only discipline providers if there is effective competition.[288] In order for disclosure to have a disciplining effect, customers need to realize that the network provider is discriminating against an application they want to use. They need to be able to switch to another provider that meets their needs and does not impose a similar restriction, and they need to be able to do so at low costs. Even if there is competition in the market for Internet access services, these conditions will often not be met, making this rule an ineffective safeguard against discrimination.[289]

First, even with disclosure, users' decision to switch will suffer from incomplete knowledge, cognitive limitations and cognitive biases. Users may not realize that their network provider is interfering with their application.[290] An application's bad performance may have many reasons (e.g., bad application design, insufficient server capacity, network congestion, problems on the network of another Internet service provider), and network provider interference will not necessarily be the first explanation that comes to mind.[291] Even if users consider that possibility, many will lack the expertise to investigate the cause of the bad performance.[292] While mandatory disclosure of discriminatory practices is intended to address this problem, experience with disclosure requirements in other contexts shows that disclosure is usually less effective at informing consumers than would be necessary for disclosure to have the intended effect.[293] Consumers often do not read disclosures, and in many cases, those who read them do not understand them.[294] For those who read and understand the disclosure, knowing which

[288] The following discussion focuses on the merits of a non-discrimination rule that does not impose any substantive limits on network providers' ability to engage in discriminatory conduct and relies solely on disclosure to discipline providers. It does not focus on the merits of mandating disclosure as a complement to substantive regulation.

[289] Relative to markets in which Internet service providers do not face any competitors, competition in the market for Internet services may even increase Internet service providers' incentives to block or discriminate. See generally van Schewick (2010d), pp. 255-259 and, regarding incentives to engage in discriminatory traffic management, Cooper (2013a); Cooper (2013b), chapters 5, 7 (both based on a case study of broadband traffic management in the UK).

[290] It is well established in the economics literature that customers' having imperfect information can provide market power to an economic actor who faces competition in the primary market by enabling the actor to impose restrictions in a complementary market that it would not be able to sustain if the primary market was perfectly competitive. See, e.g., Craswell (1982); Bar-Gill (2006).

[291] van Schewick (2010d), pp. 260-261; van Schewick (2007), pp. 376-377.

[292] For example, while user complaints about problems with BitTorrent on Comcast's network had been circulating for months, the exact method of interference was investigated and documented by Robb Topolski, a Comcast subscriber and network engineer, and later confirmed by the Associated Press and the Electronic Frontier Foundation, who had independently run their own tests upon learning of Topolski's research. See, e.g., Eckersley, von Lohmann & Schoen (2007), pp. 1-2. Since the Comcast incident, developers have created a number of tools that allow users to test their Internet connection for various signs of network provider interference. Different tools require different levels of expertise. For a list of measurement tools, see Electronic Frontier Foundation (2008). To help foster the creation of tools that consumers can use to monitor their network connections, the Federal Communications Commission in January 2011 announced a challenge to software developers and researchers "to produce research and create apps that empower consumers to monitor and protect Internet openness." The winners were announced in August 2011. Challenge.gov (2011).

[293] See, e.g., Latin (1994) (product warnings); Edwards (2005) (truth-in-lending); Cate (2006) (privacy); Ripken (2006) (securities regulation); Ben-Shahar & Schneider (2011), pp. 665-679, 704-729 (reviewing the experience with disclosure in a variety of contexts).

[294] See, e.g., Cate (2006), pp. 360-363 (citing studies on privacy notices); Edwards (2005), pp. 229-233 (citing studies on disclosures mandated by the truth-in-lending act); Ben-Shahar & Schneider (2011), pp. 666, 668-669, 671-679,

practices their network provider engages in will not necessarily allow them to make an informed decision. Many users lack the technical expertise to understand how the disclosed practices will affect them. This problem will be particularly pronounced with respect to discriminatory network management practices. Even if users understand how the practice impacts the applications they currently use, they are ill-positioned to assess the social, cultural or political consequences of the disclosed practice, its impact on future application providers' incentives to innovate or its implications for the Internet's ability to support future applications that have not yet been developed. For example, a user who believes that peer-to-peer file-sharing applications like BitTorrent are primarily used for illegal file-sharing and who does not engage in illegal file-sharing himself will not feel burdened by a network management practice that targets peer-to-peer file-sharing applications.[295] Most likely, he will not know about the various economic, technical, social, cultural and political implications of allowing this practice that were discussed above. Since they do not know the full costs of the practice, users will underestimate the benefits of switching. Moreover, many of the benefits of disciplining providers engaged in discriminatory practices by switching to another provider (e.g., more and better future applications) are in the future and uncertain, so users give them less weight than would be justified. Finally, users make the decision to switch based on an assessment of the private costs and benefits associated with switching. Since users bear the full costs of switching, but do not internalize all the social benefits of the decision to switch, they will switch less often than would be in the public interest.

Second, disclosure cannot discipline providers if there is no comparable provider to switch to who does not interfere with the applications customers want to use. Thus, the effectiveness of disclosure depends at least in part on the level of competition in the market for Internet access services. In the US, this is a real problem (see also *Box 13: Competitiveness of the Market for Broadband Internet Access in the US* below).[296,297] According to the FCC's National Broadband Plan, which was published in 2010, 78% of housing units in the US are in areas served by two wireline, facilities-based Internet access providers, while 13% are in areas

709-718 (citing studies on disclosures in a variety of contexts); Calo (2012), pp. 1050-1055 (summarizing literature on disclosure from a variety of contexts).

[295] On this and the following, see the discussion in the text surrounding footnotes 245 to 272 above.

[296] Opponents of network neutrality regulation usually have a more optimistic view of the actual amount of competition in the US than the one taken by this paper. For an example, see, e.g., Becker, Carlton & Sider (2010), pp. 502-506. In particular, they use older FCC data based on Internet service availability by zip codes (Becker, Carlton & Sider (2010), p. 503), which overstates the amount of competition (United States Government Accountability Office (2006), p. 18), treat mobile broadband Internet service as a substitute (Becker, Carlton & Sider (2010), p. 504-505) instead of treating it as a complement for wireline Internet services (Free Press (2009b), pp. 42-43, 104-105; Free Press (2009b), pp. 40-44; Free Press (2009a), p. 46 Fn. 109; Free Press (2010), pp. 45-47; Federal Communications Commission (2010b), pp. 40-41), ignore or downplay the impact of switching costs (Becker, Carlton & Sider (2010), p. 503), bundling and differentiation in the market for Internet services on the effectiveness of competition, and do not reflect the more recent decisions by AT&T and Verizon to stop expanding their fiber offerings (Becker, Carlton & Sider (2010), p. 504).

[297] The market for wireline broadband services in Europe is more competitive. See the text surrounding footnote 302 below.

where only one such provider offers service.[298] This market structure has been characterized as "duopoly +/-".[299] While a duopoly is often better than a monopoly, duopolists enjoy a degree of market power that enables them to impose restrictions on their customers that they would not be able to impose in a competitive market.[300] Mobile Internet users in the US have somewhat more options: 77% of the population lives in census tracts with three or more 3G mobile providers, 12% in areas with two providers, and 9% in areas with one.[301] In the EU, consumers usually have more providers of fixed wireline broadband service to choose from, since the regulatory framework allows unaffiliated Internet service providers to offer their services over the incumbent's network infrastructure.[302]

Box 13

Competitiveness of the Market for Broadband Internet Access in the US

Opponents of network neutrality regulation usually have a more optimistic view of the actual amount of competition in the US than the one taken by this paper.[303] In particular, they use older FCC data based on Internet service availability by zip codes,[304] which overstates the amount of competition,[305] treat mobile broadband Internet service as a substitute instead of treating it as a complement for wireline Internet services,[306] ignore or downplay the impact of switching costs,[307] bundling and differentiation in the market for Internet services on the effectiveness of competition, and do not account for the more recent decisions by AT&T and Verizon to stop expanding their fiber offerings.[308]

[298] 4% of housing units are in areas that are served by three wireline, facilities-based broadband Internet access providers (usually a DSL or fiber provider, a cable company and a cable over-builder). 5% are in areas with no wireline provider. Federal Communications Commission (2010b), p. 37. The data is based on the FCC's Form 477 data. This form counts the number of providers who offer service to at least one subscriber in a certain geographic area, without checking whether different providers offer service in overlapping geographical areas. While the National Broadband Plan tries to correct for this deficiency (Federal Communications Commission (2010b), p. 62, footnote 6), the data is likely to overstate the amount of competition to individual households. See United States Government Accountability Office (2006), p. 18.

[299] See, e.g., Farrell (2006), p. 202.

[300] See, e.g., Farrell (2006), pp. 202-205.

[301] Federal Communications Commission (2010b), pp. 39-40 (noting that these numbers "likely overstate the coverage experienced by consumers, since American Roamer [the source of the data] reports *advertised* coverage as reported by many carriers who all use different definitions of coverage"). On the market structure for mobile broadband in Europe, see, e.g. Marcus (2008), p. 35.

[302] See, e.g., European Commission (2007), pp. 18-47 (discussing market data and the existing regulatory framework); Marcus, et al. (2011), pp. 49-50; Cave & Crocioni (2011), p. 58.

[303] For an example, see, e.g., Becker, Carlton & Sider (2010), pp. 502-506.

[304] See, e.g., Becker, Carlton & Sider (2010), p. 503.

[305] United States Government Accountability Office (2006), p. 18.

[306] See, e.g., Becker, Carlton & Sider (2010), p. 504-505 (treating mobile broadband Internet service as a substitute for wireline Internet service) vs. Free Press (2009b), pp. 42-43, 104-105; Free Press (2009b), pp. 40-44; Free Press (2009a), p. 46 Fn. 109; Free Press (2010), pp. 45-47; Federal Communications Commission (2010b), pp. 40-41 (all arguing that mobile broadband Internet service is currently a complement to wireline Internet service).

[307] See, e.g., Becker, Carlton & Sider (2010), p. 503.

[308] See, e.g., Becker, Carlton & Sider (2010), p. 504.

Focusing solely on the number of providers, however, will often overestimate the number of viable alternatives available to a consumer who is willing to switch in response to discriminatory conduct. The Internet service offerings of various providers differ substantially in price, performance, and other characteristics on which providers compete.[309] As a result, even if there is another provider, switching in response to the discrimination may require a customer to switch from her most preferred Internet access offering to another offering that may meet fewer of her needs, creating an ongoing cost that will reduce the customer's willingness to switch. In the worst case, the other providers do not meet the needs of the customer at all, making it impossible for her to switch. For example, cellular providers compete on many factors such as price, coverage, devices, roaming agreements, services and, more recently, bandwidth usage caps on data plans.[310] If the other providers that do not discriminate against the application do not offer the coverage a customer needs, switching is not a realistic option. Similarly, cable networks that have been upgraded to DOCSIS 3.0 and networks offering fiber-to-the-premises are able to offer peak download speeds of more than 50 Mbps. By contrast, the peak download speeds feasible on networks offering fiber-to-the-node or on traditional DSL networks are significantly lower.[311] In the US, cable providers have generally upgraded their networks to DOCSIS 3.0, while Verizon and AT&T have stopped expanding their fiber offerings.[312] As a result, analysts expect that within a few years, 75% of the population will live in areas that will have only one service provider – the cable provider – that can offer peak download speeds of more than 50 Mbps. Only 15% of the population will likely have access to two such providers.[313] Thus, for most users interested in the highest available peak download speeds, switching providers in response to discriminatory conduct will not be a viable option.[314]

The trend toward bundling differentiates the market further, giving providers additional market power.[315] Cellular providers bundle voice, text messaging and mobile Internet access service. Wireline providers bundle telephony, television and wireline Internet access. Cable customers may not think of the digital or satellite television service offered by phone networks as a perfect substitute for their cable television; on the other hand, customers of a conventional telephony provider may not trust the digital telephony offered by cable companies.[316] Though it is possible to switch only the Internet service and keep the other offerings, this will significantly reduce the bundle discount. The problem is exacerbated if the network provider offers exclusive

[309] The following discussion draws in part on van Schewick (2010d), pp. 262.

[310] Federal Communications Commission (2010b), pp. 39-40.

Federal Communications Commission (2010b), p. 42; Crawford (2011), pp. 247-248.

[312] Crawford (2011), pp. 246-249

[313] Federal Communications Commission (2010b), p. 42.

[314] Susan Crawford has called this the "looming cable monopoly." Crawford (2010).

[315] The following discussion draws in part on van Schewick (2010d), pp. 263.

[316] In a survey of broadband users in the US, the FCC found that 39% of broadband service customers with a choice of more than one broadband provider "said that having to change their current bundle of Internet, TV, and phone service was a *major* reason for keeping service. Federal Communications Commission (2010a), p. 3.

content or exclusive devices that are valuable to the customer.[317] For example, while AT&T was the exclusive provider of the iPhone, AT&T Wireless customers may have hesitated to switch to another cellular provider that does not offer or support the iPhone.[318] Thus, product differentiation in the market for Internet services and in the market for wireline or cellular bundles makes switching to a different provider that meets fewer of their needs less attractive to customers, and gives network providers an additional degree of market power over their Internet service customers which allows them to impose restrictions they would not be able to impose in a perfectly competitive market.[319]

Even if there is more than one provider that can meet a user's needs, switching is not an option if all providers in this group engage in the discriminatory conduct.[320] For example, for several years, all mobile providers in France and Germany contractually banned the use of Internet telephony applications over mobile Internet connections. Similarly, the CRTC's review of the network management practices of Internet access service providers in Canada, where users have considerably more options for Internet access than in the US, showed that many providers were engaging in discriminatory traffic management practices that targeted peer-to-peer file-sharing applications.[321] Moreover, once discrimination is generally allowed as long as it is disclosed, different providers may discriminate against different combinations of applications, making it difficult to find a provider that meets the customer's needs and does not interfere with any of the applications the customer wants to use.

Third, the market for Internet services is characterized by significant switching costs that reduce consumers' willingness to switch and limit the effectiveness of competition.[322] Switching costs are the costs a customer incurs when switching to a competitor.[323] Switching costs make consumers' demand less elastic, enabling a provider to charge a higher price.[324] They also allow a provider to impose other restrictions that it could not impose in a perfectly competitive market. Whether these costs will prevent a customer from switching depends on the value the customer places on the excluded application and on the magnitude of the switching costs. Thus,

[317] For example, an empirical study of competition between cable television and direct broadcast satellite (DBS) multi-channel services showed that while customers generally tend to switch from cable to DBS when the quality-adjusted price of cable increases substantially, the availability of regional sports channels reduced DBS penetration, either because it raised consumers' switching costs or because it increased product differentiation between the two types of services (Wise & Duwadi (2005), pp. 695, 699–700).

[318] In September 2011, Sprint's CEO Dan Hesse noted that the fact that Sprint wasn't offering the iPhone was "the No. 1 reason customers leave or switch" (Lublin & Ante (2011)).

[319] That product differentiation may provide sellers with some degree of market power is well established in the literature (Carlton & Perloff (2005), pp. 203–205).

[320] van Schewick (2010d), pp. 259-260.

[321] For a summary of Internet service providers' responses in that proceeding, see Parsons (2009).

[322] The following six paragraphs draw in part on van Schewick (2010d), pp. 261-264. For an attempt to calculate the costs of switching broadband providers in France, see Krafft & Salies (2008); Krafft & Salies (2009). See also Cullen & Shcherbakov (2010) (estimating the explicit and implicit switching costs in the US wireless industry at approximately $230).

[323] For an in-depth overview of the economic literature on switching costs, see Farrell & Klemperer (2007). For a treatment of switching costs in the context of information goods, see Shapiro & Varian (1998), chapters 5 and 6.

[324] E.g., Varian (1999), pp. 604–605; Hausman, Sidak & Singer (2001), p. 164.

discrimination against popular applications like Google or Facebook that users view as essential will be more likely to motivate users to switch than discrimination against a newly launched application.

Switching costs in the market for Internet services are substantial. Consider first the obvious financial expenses that may be associated with switching providers. A customer who cancels a long-term contract with his provider before the end of the term will be charged an early termination fee.[325] When switching from a broadband-over-cable service to a digital-subscriber-line (DSL) service, a consumer will be charged for installation and will have to buy a DSL modem and other new equipment.[326] If (as is common in the United States) the Internet service is bundled with television and telephony, cancellation of the Internet service portion of the bundle may result in a loss, or a partial loss, of the bundle discount, and the loss of that discount may then be a significant ongoing financial cost for the consumer.[327]

Further, switching providers may require a customer to invest a significant amount of time and effort. She will have to search for and compare alternative offerings to choose a new provider. She will have to open an account with the new provider and close her account with her present provider.[328] If she cannot install the access hardware and software herself (which takes time and expertise), she must stay at home for the installation.[329] A customer who has been using an e-mail address offered by the network provider will have to notify various people of her new e-mail address, perhaps have new stationery and business cards printed, update her résumé and her website, and bear the risk of missing e-mail messages sent to the old

[325] For example, HearUsNow.org, a project of the Consumers Union, found that a number of the top broadband providers in the United States charge early-termination fees. For example, at the time of the survey in March 2007, Qwest charged a $200 early-termination fee on a two-year contract for high-speed Internet service, EarthLink charged a $149 early-termination fee on a one-year contract for DSL service, and AT&T (including SBC and BellSouth) charged a $99 early-termination fee (Consumer's Union (2007); Dunbar (2007)). In a survey of broadband users in the US, the FCC found that 32% of broadband service customers with a choice of more than one broadband provider "said paying termination fees to their current ISP was a *major* reason for keeping service." Federal Communications Commission (2010a), p. 3.

[326] In a survey of broadband users in the US, the FCC found that 50% of broadband service customers with a choice of more than one broadband provider "said paying set-up or installation fees were major factors in keeping service." Federal Communications Commission (2010a), p. 3.

[327] The customer may switch his whole bundle to the new provider, but that creates other problems, for example by making the decision to switch more complex, or by resulting in the loss of the preferred service offering, for example in television or telephony. In a survey of broadband users in the US, the FCC found that 39% of broadband service customers with a choice of more than one broadband provider "said that having to change their current bundle of Internet, TV, and phone service was a *major* reason for keeping service. Federal Communications Commission (2010a), p. 3.

[328] Providers have considerable influence over this cost. For example, in 2005, AOL paid $1.25 million in fines as part of a settlement with the state of New York, because AOL's customer service representatives were incentivized to dissuade customers from switching away from America Online, "by either making the cancellation process so painful for the customers that they could not bear to continue, or by simply ignoring their requests" (Stross (2005)).

[329] In a survey of broadband users in the US, the FCC found that 43% of broadband service customers with a choice of more than one broadband provider " said dealing with the hassle of getting new service installed was a *major* reason they have kept service." Federal Communications Commission (2010a), p. 3.

address.[330] The precise cost of switching e-mail addresses is difficult to measure, but anecdotal evidence indicates that customers view it as substantial. The *New York Times* reported in 2005 that AOL had about 5 million customers who paid $14.95 per month in order to keep using an AOL e-mail address even though they had switched to another broadband-access provider and paid Internet service fees to the new provider.[331] Medium and large businesses who switch Internet service providers will often need to renumber their networks, which is a costly, time-consuming, tedious and error-prone process.[332]

The exact costs of switching depend on the circumstances. Some customers may use provider-independent e-mail services, such as Hotmail or Gmail; others may not subscribe to a bundle at all. Some customers are not subject to a long-term contract or their contract does not include early termination fees. Sometimes, the new provider may waive the installation fee. In countries with open-access regulation where regulation allows independent Internet service providers to offer their services over other providers' networks, customers may be able to switch to another provider that offers its services over the same physical network; that removes the need to buy new equipment. Also, regulators may adopt policies to reduce switching costs. For example, the EU Universal Service Directive allows Internet service customers to switch providers in response to a change in disclosed discriminatory practices without incurring early termination fees.[333]

Thus, a particular Internet customer may face any combination of the switching costs discussed above. Every customer, however, must go through the process of searching for and choosing an alternative provider and installing and setting up the access software. These hurdles alone may deter switching. Moreover, empirical studies show that the decision to switch depends on the perceived costs of switching, which are not necessarily equivalent to the actual costs. Studies of the UK's market for long-distance telephone service have shown that providers were significantly more likely to retain dissatisfied customers who perceived the switching costs as high than dissatisfied customers who perceived them as low.[334] According to studies of the long-distance and credit-card industries, the perceived costs of switching are significantly increased if the product is perceived as complex, if it has a large number of features, or if it is bundled with other products.[335] This suggests that customers in the market for Internet access

[330] On the use of provider-specific e-mail addresses as a way to increase switching costs in Internet services, see Shapiro & Varian (1998), pp. 109–110. In other telecommunications markets such as wireline telephony and mobile telephony, regulation often requires providers to provide number portability, i.e., to enable a customer to keep its phone number when he switches providers. In 2007, the FCC asked for comments on a petition (Mortenson (2007)) to require e-mail providers to forward e-mail to a new e-mail address for a limited time (Federal Communications Commission (2007)), but did not take any further steps in this proceeding.

[331] Stross (2005). In a survey of broadband users in the US, the FCC found that 34% of broadband service customers with a choice of more than one broadband provider "said having to give up their current email address from their ISP was a major reason for not changing service." Federal Communications Commission (2010a), p. 3.

[332] Carpenter & Rekhter (1996); Carpenter, Atkinson & Flinck (2010).

[333] Article 20, Section 2. Universal Service Directive.

[334] Ranaweera & Prabhu (2003).

[335] Burnham, Frels & Mahajan (2003).

services, where services are viewed as complex, are characterized by many features and are often sold as part of a bundle, will perceive switching costs as high.

Finally, research in behavioral economics indicates that even very small switching costs may prevent customers from switching. Individuals exhibit a "status quo bias"—they are much more likely to keep what they already have than rational-choice theory would predict.[336] For example, this bias is exploited by free trials that automatically convert to a paid subscription at the end of the trial period unless the customer calls or writes to prevent this.[337] If, however, the costs of placing a call or writing a letter are sufficient to prevent people from acting, the significantly higher actual (or perceived) costs of switching Internet service providers may prevent many Internet service customers from switching providers, even if their Internet service provider excludes applications or content they would like to use.

In sum, participants in the network neutrality debate often assume that the viability of disclosure rules as a substitute for substantive regulation solely depends on the amount of competition in the market for Internet access services. After all, if there is no competition, there will be no other providers that consumers can switch to in response to discriminatory conduct, making it impossible for them to discipline providers. Based on this reasoning, participants in the debate often assume that mandatory disclosure alone will be sufficient to discipline wireline providers in Europe or in countries like Canada, where the market for wireline Internet access is generally more competitive than in the US.[338] Similar arguments are often made for mobile Internet access, where users often have a choice between three or more competitors.[339]

These arguments fail to recognize that the market for Internet services is characterized by a number of factors – incomplete customer information, product differentiation in the market for Internet access and for wireline and wireless bundles, and switching costs – that limit the effectiveness of competition and reduce consumers' willingness to switch. Rules that require network providers to disclose whether and how they interfere with applications and content on their networks reduce the problem of incomplete customer information, though only to some extent. They do not remove any of the other problems. As a result, they still leave the network provider with a substantial degree of market power over its customers, enabling it to restrict

[336] Status quo bias seems to result from a number of factors. For example, contrary to rational-choice theory, consumers often take past sunk costs into account when making consumption decisions (Samuelson & Zeckhauser (1988), pp. 37–38). Choosing one option and rejecting the other also creates cognitive dissonance, which is reduced by subsequent rationalization that the chosen option is more desirable than it was *ex ante* (Brehm (1956), p. 389). Finally, people tend to regret bad outcomes that are a result of their own action more than bad outcomes that are the result of their inaction, which again leads to a bias for doing nothing (Kahneman & Tversky (1984), pp. 343–344).

[337] Trial subscriptions with a low introductory price that automatically convert to a higher price, or other contracts with automatic renewal also exploit the cognitive bias that people tend to overestimate their future willingness to incur the then immediate costs of switching (or terminating the contract) in order to reap the future benefits (i.e., the savings) resulting from switching (or terminating the contract) (DellaVigna & Malmendier (2004), pp. 381–393).

[338] European Commission (2007), pp. 91-92; Baumol, et al. (2007); Marcus, Wernick & Carter (2008), pp. V-VI, 56-57. See also Faulhaber (2009), pp. 762-766; Faulhaber & Farber (2010b), pp. 315-316 (arguing in favor of mandatory disclosure as a substitute for substantive network neutrality rules based on a more favorable view of the amount of competition in the US).

[339] Faulhaber & Farber (2010a).

some applications and content on its network without losing too many Internet service customers.[340] They also do not affect the cognitive biases, cognitive limitations and externality problems that lead users to underestimate the benefits of switching providers compared to what would be in the public interest. Thus, even if there is competition in the market for Internet access services, disclosure cannot replace substantive regulation as a tool to discipline providers.[341]

The experience in Europe, Canada and the market for mobile Internet services in the US supports this view.

The markets for wireline Internet service in Europe and Canada are considerably more competitive than the market for wireline, fixed Internet services in the US.[342] The European legal framework does not prohibit restrictions on the end users' use of applications or services, but requires Internet access service providers to disclose them. Still, as the results of an investigation by the Body of European Regulators for Electronic Communications (BEREC) show, many Internet service customers in the European Union are subject to restrictions on their fixed or mobile Internet services.[343] A recent study showed widespread discriminatory network management in the UK.[344] In Canada, the 2009 investigation of the Canadian Regulatory Agency CRTC into Internet service providers' network management practices showed that at the time, many Canadian providers were singling out peer-to-peer file-sharing applications for special treatment, throttling the bandwidth available to them or interfering with these applications in other ways.[345]

Under the FCC's Open Internet Order, providers of mobile Internet services in the US were subject to limited restrictions on their ability to block applications and were free to discriminate, but were required to disclose, among other things, blocking of or discrimination against applications.[346] Since the adoption of the Open Internet Order, wireless carriers have engaged in various forms of discriminatory conduct, even though the market for mobile Internet services in the US is considerably more competitive than the market for wireline Internet

[340] van Schewick (2010d), p. 264; van Schewick (2007), pp. 374-377;.

[341] Relative to markets in which Internet service providers do not face any competitors, competition in the market for Internet services may even increase Internet service providers' incentives to block or discriminate. See the references cited in footnote 289 above.

[342] See footnotes 296 to 300 and accompanying text and *Box 13: Competitiveness of the Market for Broadband Internet Access in the US* above (describing US) and footnote 302 and accompanying text above (describing Europe).

[343] Body of European Regulators for Electronic Communications (2012a).

[344] Cooper (2013a); Cooper (2013b), Chapters 5, 7.

[345] For an overview of Canadian providers network management practices as disclosed during the proceeding, see Parsons (2009). Since then, most of the larger Canadian Internet service providers, most recently Bell Canada and Bell Aliant, have changed their practices in response to the regulations regarding network management that the CRTC adopted following its investigation. In January 2012, Rogers remained the only larger Canadian provider that was still engaging in discriminatory network management. Schmidt (2012); Geist (2011).

[346] 47 C.F.R. §8.3, §8.5, §8.7 and Federal Communications Commission (2010c), para 56 (describing the obligation to disclose "application-specific behavior" under 47 C.F.R. §8.3).

services.[347] Examples are Verizon Wireless' conduct towards tethering applications,[348] Verizon Wireless', AT&T's and T-Mobile's actions towards Google Wallet,[349] and AT&T's actions towards Face Time.[350] These examples suggest that at least in the market for wireline Internet service in Europe and Canada and in the market for mobile Internet services in the US, competition does not prevent Internet service providers from interfering with applications, content or services on their networks, even if, as in the US and in the European Union, network providers are required to disclose any discriminatory conduct that occurs.[351]

While mandatory disclosure alone does not sufficiently protect against discriminatory conduct, it serves many other valuable functions. Thus, it is an important complement to substantive nondiscrimination rules (see *Box 14: The Benefits of Disclosure Rules*).[352]

Box 14
The Benefits of Disclosure Rules

While mandatory disclosure alone does not sufficiently protect against discriminatory conduct, it serves many other valuable functions. Thus, it is an important complement to substantive nondiscrimination rules.

Disclosure improves competition by providing customers with information that can help them make informed decisions when choosing providers. Disclosure of traffic management practices also enables competitors to differentiate themselves along new dimensions. Today, network providers in the US compete based on maximum upload and download speed and price. If, however, customers are unable to note the differences between the offerings along other dimensions (e.g. how oversubscribed is the network, how often is traffic management used, how is traffic prioritized), they cannot take these factors into account when making a decision, and network providers do not have an incentive to compete on these factors. Thus, disclosing these characteristics along with more detailed performance measures would not only help consumers

(continued)

[347] See footnotes 296 to 301 and accompanying text above.

[348] van Schewick (2011b).

[349] van Schewick (2011a).

[350] Ziegler (2012); Kang (2012).

[351] One could argue that the existence of restricted offerings is less problematic if there are unrestricted offerings available that users can switch to. As I have explained elsewhere, this argument is not correct. The restricted offerings harm users and reduce application innovation, even if unrestricted offerings are available. See van Schewick (2012a), pp. 22-23.

[352] The following paragraph partly draws on van Schewick (2008a), pp. 1-2. Network neutrality proponents generally support adopting disclosure rules as a complement to substantive regulations. See, e.g., van Schewick (2008a), pp. 1-2; Center for Media Justice, et al. (2010), pp. 63-67; Free Press (2010), pp. 112-121; Open Internet Coalition (2010), pp. 86-92. In the US, network providers have generally argued against any mandatory disclosure rules, whether as a substitute or complement to substantive network neutrality regulation. See, e.g., AT&T (2010a), pp. 188-196; Verizon & Verizon Wireless (2010), p. 132.

make more informed choices, but also motivate Internet service providers to compete along these previously hidden dimensions.[353] More detailed disclosure of traffic management measures may also help alleviate congestion by enabling customers to adjust their behavior. Finally, disclosure provides visibility to regulators, competitors and industry observers and saves costs by removing the need for difficult and costly private investigations into a specific provider's network management practices. For example, in 2007, complaints about problems with BitTorrent and other peer-to-peer file-sharing applications on Comcast's network had circulated on user forums for several months. When asked by a reporter and later by the Electronic Frontier Foundation, Comcast denied that it was interfering with BitTorrent.[354] As a result, users, public interest organizations and reporters had to expend considerable technical effort to understand what Comcast was doing and trace BitTorrent's unusual behavior back to Comcast's intervention.[355]

Substantive Approaches

The second set of approaches in this group relies on substantive criteria to specify in advance which forms of differential treatment should be allowed. These approaches share a common goal: they seek to preserve the beneficial environment for application innovation and network use that the Internet's original architecture created in the past.

The Internet's original architecture was based on the layering principle and on the broad version of the end-to-end arguments. As a consequence of that design, the Internet was application-blind – it was unable to distinguish among the applications on the network – and, as a result, it was unable to make distinctions among data packets based on this information.[356]

As I have explained in detail elsewhere, this architecture created an environment for application innovation and network use that was application-agnostic, supported innovation without permission and user choice and kept the costs of application innovation low. These factors, in turn, allowed the Internet to foster application innovation, improve democratic discourse, facilitate political organization and action, and create a decentralized environment for cultural and political interaction in which anybody can participate.[357]

[353] See, e.g., St. Arnaud (2009), paras 26-28, 40-42, 49-51 (arguing in favor of requiring network providers to disclose oversubscription ratios); Lennett (2009), pp. 140-142 (same); IEEE-USA Committee on Communications Policy (2010) (arguing in favor of standardization and disclosure of more detailed performance measures, including metrics for bandwidth, packet loss, latency, jitter and availability), pp. 12-18, 24; IEEE-USA (2010), pp. 1-3 (same).

[354] Reardon (2007); Schoen (2007). This sentence and the next sentence are adapted from van Schewick (2010d), p. 261.

[355] Comcast's method of interfering with BitTorrent was first investigated by Comcast subscriber and network engineer Robb Topolski, who detected the spoofed RST packets that Comcast was using to reset BitTorrent connections. Upon learning of Topolski's research, the Associated Press and the Electronic Frontier Foundation independently ran their own tests and documented the practice. Svensson (2007); Eckersley, von Lohmann & Schoen (2007). .

[356] See footnotes 2 and 55 above.

[357] On these factors and their economic, social, cultural and political impact, see footnotes 50 to 59 and accompanying text, *Box 4: Application-Agnostic v. Application-Blind* and *Box 3: The Importance of User Choice* above.

Today, technologies such as Deep Packet Inspection have removed the application-blindness of the network. They allow network providers to identify the applications and content on their networks and to control their execution.[358]

In response, the non-discrimination rules in this section try to preserve through law the environment for application innovation and network use – an environment characterized by application-agnosticism, user choice, innovation without permission and low costs of application innovation – that the Internet's original architecture created by virtue of its architectural design. Put differently, the rules in this section seek to preserve the Internet's ability to function as a general-purpose platform over which applications, content, services and uses compete on a level playing field, with users choosing which applications become successful and how the network can be used. They differ, however, in their assessment of which behavior needs to be banned in order to realize this goal (see *Table 1: Similarities and Differences Among the Approaches*).[359]

Table 1
Similarities and Differences Among the Approaches

	Discrimination among like applications	**Discrimination among like classes of applications**	**Discrimination among classes of applications that are not alike**	**Application-agnostic discrimination**
First approach	Banned	Banned	Allowed	Allowed
Second approach	Banned	Banned	Banned	Allowed

According to the first approach, discriminatory conduct distorts competition among applications or classes of applications only if it differentiates among like applications or classes of applications. In line with this assessment, the approach bans discrimination among like applications and among like classes of applications, but allows discrimination among classes of applications that are not alike. In other words, this approach requires network providers to treat like traffic alike. It also allows discrimination that is application-agnostic. (The terms "application" and "class of application," are defined in *Box 16: Terminology: "Application" and "Class of Application"* below.)

According to the second approach, which is the approach proposed by this paper, any differential treatment that is application-specific interferes with the values that network neutrality regulation is designed to protect. (Differential treatment is application-specific if it is based on application or class of application, or, put differently, if it is based on criteria that depend on an

[358] On Deep Packet Inspection (DPI) in general, see Anderson (2007). For a specific example, see Cisco Systems (2005a). On the state of DPI deployment, see Free Press (2010), pp. 141-151.

[359] As the discussion will show, only the second approach meets this goal.

application's characteristics.)[360] In line with this assessment, the approach bans application-specific discrimination, but allows application-agnostic discrimination.

This approach bans all discrimination among applications and classes of applications that is based on application-specific criteria, regardless of whether the applications or classes are alike or not. Thus, the first and the second approach differ in how they treat discrimination among classes of applications that are not alike (see *Table 1: Similarities and Differences Among the Approaches* above).

Ban Discrimination Among Like Applications or Classes of Applications, but Allow Discrimination Among Classes of Applications That Are Not Alike and Application-Agnostic Discrimination

The first approach prohibits only discrimination among like applications or classes of applications, but allows discrimination among classes of applications that are not alike and application-agnostic discrimination.[361] (Again, I use "applications" as shorthand for "applications, content, services, and uses.") Thus, the approach requires network providers to treat like traffic alike. This requirement is often called "like treatment."[362] The non-discrimination rule in the merger conditions of the AT&T/BellSouth merger has been interpreted as requiring like treatment. (See *Box 15: Like-Treatment and the AT&T/Bell South Merger Conditions*) (The terms "application" and "class of application" are defined in *Box 16: Terminology: "Application" and "Class of Application*" below.)

This approach assumes that the criterion that distinguishes socially beneficial from socially harmful differential treatment is whether the applications or classes of applications that are being treated differently are alike or not. If they are alike, the differential treatment is socially harmful and should therefore be banned. If they are not alike, the differential treatment is socially beneficial (or at least not socially harmful) and should therefore be allowed.

Although the outcome of the rule turns on whether applications are alike, proposals in this category usually do not specify how network providers or regulators should make this

[360] For a more detailed discussion of the terms "application-specific", "discrimination based on application" and "discrimination based on class of application," see footnotes 381, 437 to 439, 444 to 445 and accompanying text and the discussion in footnote 439 below. See also *Box 22: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination* below.

[361] Application-agnostic discrimination is discrimination that is based on criteria whose application does not depend on an application's characteristics. The rationale for allowing application-agnostic discrimination is set out in Section "Ban Application-Specific Discrimination, But Allow Application-Agnostic Discrimination" below. See also *Box 22: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination* below.

[362] On like treatment, see Wu (2006c), pp. 42-43. Critical of allowing like treatment: Riley & Topolski (2009); Center for Democracy & Technology (2010), pp. 29-30, 40 (arguing against allowing like treatment in the non-discrimination rule and as reasonable network management); Free Press (2010), pp. 101-104 (discussing like treatment in the context of the reasonable network management exception). In favor of allowing like treatment, see, e.g., AT&T (2010a), pp. 187-188; Covad Communications Company (2010), p. 7-8; Cox Communications (2010), pp. 23-30. See also Jordan & Ghosh (2010), pp. 12:10-12:11, 12:15, 12:19-12:20 (classifying like treatment as a "borderline traffic management practice that could be used for a limited period of time if properly disclosed in the user contract." ibid., p. 12:15).

determination.[363] Thus, the rule leaves a key term undefined. At the same time, the term "like" can be interpreted in a number of ways. The resulting ambiguity is at the heart of many problems with this rule.[364]

Box 15

Like Treatment and the AT&T/Bell South Merger Conditions

The non-discrimination rule in the merger conditions in the AT&T/BellSouth merger has been interpreted as requiring like treatment: It prohibited AT&T/Bell South to "privilege[], degrade[] or prioritize[] any packet transmitted over AT&T/BellSouth's wireline broadband Internet access service based on its source, ownership or destination."[365] According to Tim Wu, "[w]hile the agreement does not use the word discrimination, it effectively bars discrimination on the basis of source, ownership, or destination. It forbids AT&T from, for example, selling Yahoo or CNN priority access to its customers over its broadband networks, and favoring those content sources over unaffiliated blogs or search engines. […] Interestingly, the agreement does not prevent AT&T from treating different media carried on the internet differently, so long as the carrier does not discriminate between who is providing the content. […] In short, AT&T must treat like traffic alike."[366] Not all observers agree with this characterization of the agreement, though.[367]

[363] See the rules discussed in *Box 13: Like Treatment and the AT&T/Bell South Merger Conditions* and in footnote 380 below. For an exception in the context of the reasonable network management exception, see the Google/Verizon Legislative Framework Proposal which included an exception for reasonable network management that allowed network providers "to prioritize general classes or types of Internet traffic, based on latency." (Google & Verizon (2010), p. 1.) Discussions of like treatment usually do not discuss how a network provider or a regulator applying the rule should decide which applications are alike, either (see, e.g., Wu (2006a); Wu (2006c), pp. 42-43). As the literature on like treatment in the context of international trade law shows, deciding whether two applications are alike is a complex problem. (See, e.g., Quin (2005); Van den Bossche (2008), Chapter 4.)

[364] See footnotes (and accompanying text) 404 to 423 and the Section "Certainty and Costs of Regulation" below.

[365] AT&T (2006), p.8: "AT&T/BellSouth also commits that it will maintain a neutral network and neutral routing in its wireline broadband Internet access service. This commitment shall be satisfied by AT&T/BellSouth's agreement not to provide or to sell to Internet content, application, or service providers, including those affiliated with AT&T/BellSouth, any service that privileges, degrades or prioritizes any packet transmitted over AT&T/BellSouth's wireline broadband Internet access service based on its source, ownership or destination."

[366] Wu (2006a).

[367] National Cable & Telecommunications Association (2007), pp. 17-19 (disagreeing with Wu's interpretation of the merger conditions as allowing like treatment.

Box 16
Terminology: "Application" and "Class of Application"

In this paper, the term "application" refers to a specific instance of a specific type of application. For example, Vonage is an application, as are Skype or Google Voice – each of them is a specific instance of Internet telephony applications. Gmail is one of several e-mail applications. A "class of applications" is a group of individual applications that share some common characteristic. For example, "Internet telephony" or "Internet telephony applications" (i.e. the group of all Internet telephony applications), "latency-sensitive applications" (i.e., the group of all latency-sensitive applications), or the group of all applications that use a specific application-layer or transport layer protocol (e.g., all applications that use the BitTorrent protocol) are all classes of applications.

The non-discrimination rule described in this section does not restrict how network providers define classes of applications. It only requires that once a network provider has defined different classes of applications, it must treat like classes of applications (and, of course, all applications within a class of like applications) alike. Contrary to the terminology used in this paper, participants in the debate sometimes use the term "application" to denote an application type. For them, e-mail would be an application. By contrast, under the terminology used in this paper, e-mail is a class of applications (the group of all e-mail applications). Others use the term "class of applications" synonymously with application type. Under the terminology in this paper, the term "class of application" is broader than that. In this paper, the group of applications that have the same application type (e.g., all e-mail applications, or, shorter, e-mail) is one potential class of applications, but beyond that, any group of applications that share a common characteristic can be a "class of applications."

Banning Discrimination Among Like Applications or Classes of Applications

The first part of the rule bans discrimination among like applications or classes of applications. This prevents network providers from singling out one or more specific applications within a group of like applications (or one or more of several like groups of applications) for differential treatment. For example, Comcast could not treat video streaming from Amazon or video streaming from XFinity.com, Comcast's own Internet streaming video offering, differently from video streaming from other providers such as Netflix or YouTube.[368] Similarly, Comcast would not be allowed to count traffic from other providers' streaming video applications towards its monthly bandwidth usage cap, while exempting traffic from XFinity.com from the cap.[369,370] This would be discrimination among like applications.[371]

[368] Thus, it does not matter whether the favored application is affiliated with the network provider or not.

[369] As explained above, the non-discrimination rules discussed in this paper apply to all forms of differential treatment that make some application or classes of applications relatively more attractive, not just to differential handling of packets in the network. See footnotes 75 to 84 above and accompanying text.

[370] Comcast's Internet service offerings have a monthly usage cap of 250 GB. Data traffic from XfinityTV.com and from other streaming video providers all count towards that cap. However, data traffic generated by Comcast's XfinityTV app for the Xbox, an app that allows users to view video on demand content from Comcast on an Xbox connected to the Internet through Comcast's Internet service, does not count towards the monthly bandwidth cap, while video traffic generated by other apps (e.g., the HBO GO app) for the Xbox does count towards the cap. That behavior would violate the rule described in the text. For sources and additional discussion, see *Box 6: Differential*

With respect to Quality of Service, the ban on discrimination among like applications or classes of applications prevents network providers from offering Quality of Service exclusively to some, but not all applications within a class of like applications or only to one of several classes of applications that are alike. For example, Comcast would be prohibited from providing an enhanced type of service only to XfinityTV.com, but not to unaffiliated streaming video applications. Under network neutrality regimes that allow a network provider to charge application providers for prioritized or otherwise enhanced access to the network provider's Internet access customers,[372] the ban on discrimination among like applications or classes of applications would prohibit a network provider from selling an enhanced service exclusively to one of several video conferencing providers that are not affiliated with the network provider. Thus, the ban addresses the concern that network providers may use the selective provision of Quality of Service as a tool to distort competition among applications or classes of applications.[373]

The ban on discrimination among like applications or classes of applications is designed to prevent network providers from discriminating against specific applications within a class of like applications or against like classes of applications as a substitute for blocking them. As has been set out above, discrimination is often an attractive alternative to blocking, since it is less costly and potentially more effective.[374] Thus, in cases in which a network provider has an incentive to block an application or class of applications – e.g., to manage congestion, to block unwanted content or to give an advantage to another, competing application in a way that increases the network provider's profits –[375] it often has an incentive to reach the same result by treating the targeted applications relatively worse than others (either by treating the other applications better or the targeted applications worse), and the rule is designed to prevent this.[376]

If a network provider singles out one or more specific applications within a group of like applications (or one or more of several like groups of applications) for differential treatment, the

---

Counting of Traffic Towards the Monthly Bandwidth Cap: Comcast's Xfinity TV App for the Xbox above and *Box 18: Examples of Differential Impact: Comcast's Digital Voice service and Comcast's Xfinity TV App for the Xbox* below.

[371] In these examples, the group of like applications is the group of all streaming video applications.

[372] There is considerable debate over whether a network provider should be allowed to charge application providers who are not its Internet access customers for prioritized or otherwise enhanced access to the network provider's Internet access customers. As I explain elsewhere, network providers should be prohibited from imposing such charges (van Schewick (2010h)). See also *Box 2: Charging for Quality of Service* above and footnotes 23-24 and accompanying text above.

[373] This concern is discussed by, e.g., van Schewick (2010h), p. 3; van Schewick (2010i), p. 6; Jordan & Ghosh (2010), pp. 12:14, 12:20 and the sources cited in footnote 103 above.

[374] See the discussion in the text surrounding footnotes 91 to 99 above.

[375] On incentives to block, see, e.g., van Schewick (2008a) , pp. 5-6; van Schewick (2010d), pp. 264-266 (manage congestion); van Schewick (2010d), 266-270 (block unwanted content); van Schewick (2010d), pp. 222-264 (increase profits).

[376] This paper assumes that the case for a rule against blocking has been made. See footnote 22 and accompanying text. Banning blocking, but allowing discrimination would make the rule against blocking meaningless, so the arguments in favor of a rule against blocking justify this part of the non-discrimination rule as well.

harm to the values that network neutrality regulation is designed to protect is obvious. In this case, the differential treatment – whether it treats the targeted applications better or worse – immediately reduces the relative performance of some applications in the group, making them less attractive to users than the others. Thus, the differential treatment effectively imposes a tax on some applications in the group. Compared with an application-agnostic network, where users choose among applications without interference from network providers, this distorts users' choices among applications and, as a result, tilts the playing field in favor of some applications in the group. The differential treatment distorts competition among the applications in the group and reduces the value of the network for users by manipulating them to use applications that they would not necessarily have chosen otherwise. It also affects application innovation in various ways. In particular, the threat of discrimination reduces application developers' incentives to innovate and their ability to get funding.[377] Moreover, letting users, not network providers, pick winners and losers on the Internet is an important part of the mechanism that produces innovation under uncertainty.[378]

Allowing Discrimination Among Classes of Applications That Are Not Alike

Non-discrimination rules in this category do, however, allow network providers to differentiate among classes of applications that are not alike, as long as they do not differentiate among applications within each class. With respect to Quality of Service, they would allow network providers to offer or apply different types of service to different provider-defined classes of applications, as long as they do not discriminate among classes of applications that are alike or discriminate among like applications within a class.

Under an interpretation of the term "like" that considers applications or classes of applications to be "alike" if they have similar requirements with respect to throughput, jitter or delay, the rule would allow network providers to provide low-delay service to Internet telephony, but not to e-mail. Internet telephony is sensitive to delay, while e-mail is not, so this would be discrimination among two classes of applications that are not alike. Thus, network providers could treat Vonage, an Internet telephony application, differently from Gmail, an e-mail application, but they could not to treat Skype, another Internet telephony application, differently from Vonage, or Gmail differently from Hotmail. By contrast, this interpretation of "like" would ban providing low-delay service to online gaming, but not to Internet telephony. Online gaming applications and Internet telephony applications are both sensitive to delay, so this would be discrimination among like classes of applications.[379]

[377] For a more detailed explanation, see footnotes 229 to 231 above.

[378] For a more detailed explanation, see footnotes 46 to 54 above. See also *Box 3: The Importance of User Choice* above.

[379] Usually, neither proposals for rules requiring like treatment nor discussions of like treatment discuss how a network provider or a regulator applying the rule should decide which applications are alike (see footnote 363 and accompanying text). As the literature on like treatment in the context of international trade law shows, deciding whether two applications are alike is a complex problem. (See, e.g., Quin (2005); Van den Bossche (2008), Chapter 4.) As will be set out in more detail below, the interpretation in the text is not the only possible interpretation. This

The rule is agnostic as to who controls (or makes the decision regarding) the actual provision of the different types of service. (See *Box 17: Like Treatment and Control over the Provision of Quality of Service.*) Apart from the AT&T/BellSouth Merger conditions, several network neutrality bills introduced in Congress would have allowed Internet service providers to offer Quality of Service to applications of the same type, as long as they do not discriminate among applications of that type.[380]

Box 17

Like Treatment and Control over the Provision of Quality of Service

The rule described in the text is agnostic as to who controls (or makes the decision regarding) the actual provision of the different types of service. In a *user-controlled system*, the provider defines the classes and decides which type of service, if any, to offer to the different classes, but the user decides whether to take advantage of that possibility for applications for which a special type of service is available. For example, in the example of Shaw described in *Box 20: Defining "Like" Based on Use* below, the Canadian Internet service provider Shaw gave users the option to buy an enhanced type of service for Internet telephony applications. It did not provide the option to buy an enhanced type of service for any other class of applications. Thus, Shaw defined the class of applications (Internet telephony applications) to which it would offer the enhanced type of service, but users decided whether they wanted to actually buy that option. If they did not buy the option, their Internet telephony applications did not receive the enhanced type of service. Whether this offering complies with the rule described in the text depends on whether defining "like" based on use is an acceptable interpretation of like.[381]

In a *provider-controlled system*, the provider defines the classes, determines which class should get which Quality of Service and provides the actual service without any involvement by the user. The trial of network management practices by Cox Communications described in *Box 19: Defining "Like" Based on Application Requirements* below is an example of a provider-controlled approach. Given the concerns described in *Box 19: Defining "Like" Based on Application Requirements* below, it is questionable whether that trial treated like traffic alike, so Cox' system may not comply with the rule described in the text.

ambiguity of the term "like" is one of the key problems of this rule. See footnotes (and accompanying text) 404 to 423 below and Section "Certainty and Costs of Regulation" below.

[380] Internet Freedom and Nondiscrimination Act of 2006 (2006), §28(b): "If a broadband network provider prioritizes or offers enhanced quality of service to data of a particular type, it must prioritize or offer enhanced quality of service to all data of that type (regardless of the origin or ownership of such data) without imposing a surcharge or other consideration for such prioritization or enhanced quality of service." See also Internet Freedom Preservation Act (2007), §12(a)(5): A broadband service provider shall "only prioritize content, applications, or services accessed by a user that is made available via the Internet within the network of such broadband service provider based on the type of content, applications, or services and the level of service purchased by the user, without charge for such prioritization."

[381] The rule described in the text does not specify how to determine whether different applications or classes of applications are alike. In this case, the question whether a definition of like that treats similar uses alike is consistent with the non-discrimination rule described in text would have to be decided in future adjudications. This question is beyond the scope of this paper. See also footnotes 363 to 364 above and Sections "Application-Agnosticism and the Ambiguity of Like" and "Certainty and Costs of Regulation" below.

The decision to allow discrimination among applications or classes of applications that are not alike is based on the assumption that this kind of discrimination is socially harmless and does not threaten the values that network neutrality regulation is designed to protect. As will be set out below, this assumption is not correct. In many cases, discrimination among classes of applications hurts some classes of applications, even if the classes are not alike.

More generally, rules requiring like treatment create considerable social costs. Like treatment negatively affects several of the factors that have fostered application innovation in the past. It removes the application-agnosticism of the network and gives network providers discretion to decide which applications are alike, which allows network providers to deliberately or inadvertently distort competition among applications or classes of applications. It violates the principle of user choice, resulting in levels of Quality of Service or differential treatment that do not necessarily meet users' needs. It violates the principle of innovation without permission, reducing the chance that new applications actually get the type of service they need. Like treatment also creates high costs of regulation.

Impact of Discrimination among Classes of Applications that Are Not Alike

In some cases, discrimination among classes of application that are not alike does not harm the applications that get relatively worse treatment. For example, e-mail and Internet telephony have different requirements with respect to reliability and delay:[382] E-mail requires reliable data transfer, but is not sensitive to delay. By contrast, Internet telephony can deal with a certain amount of packet loss, but is very sensitive to delay above a certain level. As a result, it does not harm e-mail if a network provider gives low-delay service to Internet telephony, but not to e-mail, as long as the delay faced by e-mail and other best-effort traffic does not increase above a level where it negatively affects even applications that generally work well with best-effort traffic.[383] Similarly, during times of congestion, a network provider may want to prioritize applications that are time-sensitive over those that are not. An application that is not time-sensitive does not suffer if it is not prioritized during times of congestion, as long as the delay for non-time-sensitive traffic does not rise above a level that would negatively affect even non-time-sensitive traffic. In these examples, the applications that do not receive the "better" treatment are not harmed because they do not need the better treatment anyways. The differential treatment benefits the applications that get better treatment without harming any of the others, so it does not seem to interfere with competition among applications or user choice.[384]

[382] See footnote 10 above and accompanying text.

[383] How exactly the provision of an enhanced type of service to some traffic will affect the remaining best effort traffic depends on the specific mechanism used to provide the differential treatment. For example, priority queuing allows a router's high-priority queue to starve lower-priority queues for bandwidth. By contrast, fair queuing provides a guaranteed minimum share of bandwidth to the different queues. Peterson & Davie (2012), pp. 494-499.

[384] As I explain below, in reality the impact of "like treatment" on users and application providers is more complicated even in the case of these examples. For example, a network provider may deliberately or inadvertently assign an application to the wrong class. See Section "Application-Agnosticism and the Ambiguity of 'Like'" below.

These cases, however, are only a subset of the cases in which a network provider has an incentive to discriminate among classes of applications. In many cases, discrimination among classes of applications hurts some classes of applications, even if the classes are not alike.

For example, some Internet applications compete with network-provider applications that are sold separately from Internet access and do not run over the Internet-access portion of the network provider's access network. In these cases, discriminating against all applications in that class allows the network provider to favor its own offering without discriminating among applications within the class. A cellular or wireline network provider, for example, may have an incentive to reduce the performance of Internet telephony applications to protect its revenue from its own, separate telephony offering.[385] The same incentive applies to Internet video offerings that may motivate users to cancel their subscription to the network provider's traditional multi-video programming services,[386] or to Internet messaging applications that threaten cellular carriers' revenue from traditional text-messaging services.[387] In the future, this incentive will apply to every application that is offered both over the Internet portion of the network provider's access network and over the portion of the access network that is dedicated to "specialized services."[388] Since the non-discrimination rule only applies to a network provider's Internet access offering, like treatment only bans discrimination among the Internet telephony offerings (or among any other affected group of applications) that run over the Internet access portion of its network. It does not protect these applications against behavior that applies equally to all of them, but puts them at a disadvantage with respect to the network provider's offering that is sold and operated separately from Internet access. The Comcast case illustrates this problem. (See *Box 18: Examples of Differential Impact: Comcast's Digital Voice Service and Comcast's Xfinity TV App for the Xbox* below.) Thus, applications in a class can be harmed by differential treatment if that treatment puts them at a disadvantage compared to another, competing application that is outside the scope of the non-discrimination rule.

[385] See, e.g., van Schewick (2010d), pp. 240-243. See also AT&T (2009), pp. 6-7 (noting a contractual agreement between AT&T and Apple "that Apple would not take affirmative steps to enable an iPhone to use AT&T's wireless service to make VoIP calls" because "both parties required assurances that the revenues from the AT&T voice plans available to iPhone customers would not be reduced by enabling VoIP calling functionality on the iPhone." Ibid., p. 7)

[386] See, e.g., U.S. Department of Justice (2011), pp. 11, 14-20, 37-39 (citing "[m]any internal Comcast documents" showing that Comcast views online video distributors as a competitive threat to its traditional cable video distribution offerings, ibid., p.19, and describing Comcast's incentives to discriminate against unaffiliated online video providers).

[387] BBC News (2012); Chen (2012); Thomas (2012) (all discussing the reduction in mobile carriers' text messaging revenues resulting from instant messaging applications).

[388] Whether this incentive exists depends on how specialized services are regulated. See, e.g., Speta (2011).

Box 18

Examples of Differential Impact: Comcast's Digital Voice Service and Comcast's Xfinity TV App for the Xbox

Since the non-discrimination rule only applies to a network provider's Internet service offering, it does not protect a group of like applications against behavior that applies equally to all of them, but puts them at a disadvantage with respect to the network provider's offering that is sold and operated separately from Internet access.

The aftermath of the FCC's Comcast/BitTorrent order illustrates this problem. Under Comcast's new application-agnostic network management practices, when a part of the network gets close to being congested (as measured by average levels of utilization over a certain number of minutes), the traffic of users that have used a high amount of bandwidth over a certain number of minutes receives relatively less priority than the traffic of other users.[389] If the affected users are running an Internet telephony application while their traffic is treated like this, the delays resulting from the traffic management may reduce the performance of the Internet telephony application. By contrast, Comcast's own digital voice service, which is sold separately from Internet access service, is not affected by this problem. The digital voice traffic is separated from the user's Internet traffic and, therefore, not affected by any traffic management measures that are applied to that traffic. When the FCC asked Comcast to justify "the disparate treatment of its own VoIP service as compared to that offered by other VoIP providers on its network," Comcast argued that the fact that its network management practices apply to unaffiliated Internet telephony applications, but not to its own digital voice offering does not violate the FCC's Order against Comcast, since its digital voice offering is not offered over the public Internet and therefore not subject to the requirements imposed by that order.[390]

The current controversy over Comcast's Xfinity TV App for the Xbox provides another example of this phenomenon. As explained above, the Xfinity TV App allows users who subscribe to Comcast's Internet access service, Comcast's cable service and Microsoft's Xbox Live Gold subscription service to watch selected video-on-demand content from Comcast on the Xbox. Comcast's Internet service has a 250 GB monthly bandwidth cap.[391] Traffic associated with the Xfinity TV App to the Xbox does not count towards that cap, while traffic of other applications that also allow users to view on-demand video content on the Xbox (e.g., HBO Go or Netflix) does count towards the cap.[392] In general, this differential treatment of like applications would violate the non-

(continued)

[389] A user's traffic continues to receive relatively less priority until his bandwidth use has fallen below a predetermined level. For a detailed description of the system, see Bastian, et al. (2010).

[390] See Federal Communications Commission (2009a), pp.1-2; Zachem (2009), pp. 2-3. For a brief description of Comcast's network management system, see Lennett (2009), pp. 119-120.

[391] Comcast Corporation (2012a).

[392] Comcast Corporation (2012b).

discrimination rule described in this section.[393] Comcast, however, claims that contrary to the other applications, which are delivered over the public Internet and treated in compliance with the FCC's Open Internet Rules, the Xfinity TV App is provided separately from the public Internet and therefore not subject to the FCC's Open Internet Rules.[394] If this argument is correct,[395] Comcast is able to put applications that deliver online video to the Xbox over the public Internet at a disadvantage compared to Comcast's own competing online video application, even though all online video applications delivered over the public Internet are treated alike.

Moreover, applications in a class can be harmed by differential treatment even if they do not compete directly with applications in other classes that are treated more favorably. As I have explained elsewhere, network providers often have an incentive to single out specific applications or classes of applications for special treatment in order to manage bandwidth on their network.[396] For example, at the time of the Canadian investigation into Internet service providers' network management practices, many Canadian providers were singling out peer-to-peer file-sharing applications for special treatment, throttling the bandwidth available to them or interfering with these applications in other ways.[397] In the US, Comcast, RCN and, most likely, Cox for a while managed traffic on their networks by selectively interfering with BitTorrent and other peer-to-peer file-sharing applications, but not with other applications.[398] In 2009, British Telecom (BT) throttled streaming video of users subscribing to its "Up to 8 Mbps Option 1" broadband plan to 986 kilobytes/sec between 5 pm and midnight.[399] And according to Neelie Kroes, Vice President of the European Commission responsible for the Digital Agenda, new data published by BEREC, the Body of European Regulators for Electronic Communications, in June 2012, shows that around 20% of fixed Internet service providers (spread across virtually all EU member states) impose restrictions on peer-to-peer file-sharing applications during peak times. These restrictions can affect up to 95% of users in a country.[400] When a network provider singles out a class of like applications for special treatment without discriminating among

[393] See footnotes (and accompanying text) 75 to 76, 369 to 370 and *Box 6: Differential Counting of Traffic Towards the Monthly Bandwidth Cap: Comcast's Xfinity TV App for the Xbox* above.

[394] Werner (2012).

[395] Whether the Open Internet Rules apply to the Xfinity TV App is outside the scope of this paper.

[396] van Schewick (2008a), pp. 5-6; van Schewick (2010d), pp. 264-266.

[397] See Parsons (2009). Since then, most of the larger Canadian Internet service providers, most recently Bell Canada and Bell Aliant, have changed their practices in response to the regulations regarding network management that the CRTC adopted following its investigation. In January 2012, Rogers remained the only larger Canadian provider that was still engaging in discriminatory network management. Schmidt (2012); Geist (2011).

[398] Comcast Corporation (2008a); RCN Corporation (2010). Cox seems to have actively managed peer-to-peer file-sharing in 2008 as well. Schatz (2008) (citing a Cox statement that "Cox allows the use of file-sharing and peer-to-peer services for uploads and downloads, and we allow access to all legal content, but we must manage the traffic impact of peer-to-peer services, as most ISPs do for the benefit of the customer."); Dischinger, et al. (2008) (study finding evidence of BitTorrent blocking by Comcast and Cox).

[399] Cellan-Jones (2009). See also Cooper (2013a); Cooper (2013b), chapters 5,7 (documenting widespread discriminatory network management in the UK).

[400] The text draws heavily on Kroes' description of these finding in her blog (Kroes (2012)). For the detailed findings, see Body of European Regulators for Electronic Communications (2012b).

applications within the class, the resulting harm may be less apparent than in cases in which the network provider discriminates against specific applications within a class. After all, if all applications in the class are treated the same, they still compete with each other on a level playing field. Focusing only on competition among the applications within a class is too narrow. On the Internet, different uses constantly compete for users' time and attention. Differential treatment that treats a certain class of applications worse than others in a way that harms their usability or attractiveness to users (as opposed to differential treatment that does not harm the affected applications because they do not need the better treatment) imposes a tax on the developers and users of the affected application that affects user behavior and the applications' chances in the marketplace. As the co-founders of the online video company Zediva explained in a letter to the FCC: "Discriminatory network management of this type [that singles out specific applications or classes of applications in order to deal with congestion] would put the affected applications at a severe disadvantage. Companies that offer these applications and services will be less able to reach their users during times of congestion, which in turn may affect their success in the market (who wants to use an application or service that is less usable during peak time, when most people actually want to use the Internet?) and their ability to get funding – thus squashing innovation before it has had a chance to prove itself in the marketplace."[401] Differential treatment that makes a class of application less usable or attractive to users also harms users whose applications are affected by the differential treatment. It constrains their ability to use the Internet as they see fit either generally, or, when the differential treatment is used for congestion management, during peak times, when people want to use the Internet most.[402] Thus, treating classes of applications differently may harm users and applications even if the classes of applications are not alike.

Application-Agnosticism and the Ambiguity of "Like"

In addition, "like treatment" negatively affects several of the factors that have fostered application innovation in the past.[403]

In order to implement "like treatment," network providers need to identify the different applications on their network in order to decide which class they belong to and determine the appropriate form of Quality of Service or differential treatment.[404] Thus, like treatment requires network providers to treat data packets differently based on information about the applications on the network, which removes the application-agnosticism of the network. Since the concept of "like" applications is not well-defined, network providers have broad discretion when defining classes of applications or determining which class a specific application should be assigned to.

[401] Srinivasan & Gupta (2010). As the Zediva letter and conversations with entrepreneurs and investors show, this is not a theoretical concern. For another publicly documented example, see van Schewick (2008b), p. 2

[402] van Schewick (2010b).

[403] van Schewick (2010f), pp. 11-13.

[404] Center for Democracy & Technology (2010), p. 29.

This allows them to deliberately or inadvertently distort competition among applications or classes of applications.[405]

Often, there may be different options for determining which applications are "alike" and should therefore receive the same treatment. For example, one approach may focus on applications' requirements with respect to throughput, delay or jitter. (See *Box 19: Defining "Like" Based on Application Requirements* below.) Another approach may focus on whether the applications are used for similar goals, or whether they compete with each other. (See *Box 20: Defining "Like" Based on Use* below.) Depending on which option is chosen, a specific application or type of application may receive very different treatment. Thus, a network provider can put certain applications or classes of applications at a disadvantage by choosing a definition of "like" that hurts that application or class of application. Sometimes, this may happen deliberately; sometimes, it may happen inadvertently.

Box 19
Defining "Like" Based on Application Requirements

A definition of "like" could focus on whether applications have similar requirements with respect to throughput or delay. For example, the Google/Verizon Legislative Framework Proposal included an exception for reasonable network management that allowed network providers "to prioritize general classes or types of Internet traffic, based on latency."[406] In 2009, Cox Communications, a cable provider in the US, trialed a network management system that constitutes an example of such an approach. Cox divided applications into two groups: time-sensitive applications and non-time-sensitive applications. During times of congestion, the system de-prioritized applications that Cox had classified as non-time-sensitive to improve the performance of applications that Cox had classified as time-sensitive. Cox performed the classification based on "our network engineering expertise and our customers' expectations."[407] For example, web, Internet telephony, e-mail or streaming video were classified as time-sensitive, while, e.g., file access, software updates or peer-to-peer protocols were classified as non-time-sensitive.[408]

There are reasons to believe that this system would have violated the non-discrimination rule described in this section: It seems to have classified all applications that use peer-to-peer file-sharing protocols as non-time-sensitive, even though some peer-to-peer file-sharing applications (e.g., Vuze, an application that uses peer-to-peer file-sharing protocols to stream video in real time) are sensitive to delay. Treating some time-sensitive traffic as time-insensitive and, therefore, differently from other time-sensitive traffic would have violated the requirement to treat like traffic alike.[409]

[405] van Schewick (2010f), pp. 11-12.

[406] Google & Verizon (2010), p. 1.

[407] Cox Communications Inc. (2009).

[408] Cox Communications Inc. (2009); Cox Communications (2010), pp. 24-30 & Appendix A. See also footnote 417 below and Riley & Scott (2009), pp. 6-8 (criticizing Cox' approach).

[409] See footnotes 412 to 417 and accompanying text.

Box 20
Defining "Like" Based on Use

A definition of "like" could focus on whether the applications are used for similar goals. For example, in the mid-2000s, the Canadian cable provider Shaw Communications allowed its Internet service customers to add a Quality of Service enhancement option to their normal Internet service for $10 per month. This option provided enhanced Quality of Service only to Internet telephony applications, but not to other latency-sensitive applications.[410] Thus, Shaw defined the class of application to which it offered an enhanced type of service based on the use of the application ("Internet telephony"), not on applications' technical needs.

It is not clear whether such a definition of like would be allowed under the non-discrimination rule described in this section.[411]

Network providers may deliberately or inadvertently define classes in a way that hurts specific applications within a class. The Canadian Radio-television and Telecommunications Commission's (CRTC) review of the Internet traffic management practices of Internet service providers illustrates how this may happen. The proceeding showed that many Canadian Internet service providers throttled or otherwise interfered with traffic belonging to peer-to-peer file-sharing applications all day or during times of congestion.[412] They argued that this was necessary to protect the performance of real-time applications (such as applications that stream video in real time) during times of congestion.[413] This raised an interesting question: How did the network providers treat Vuze, an application that, at the time of the proceeding, used the BitTorrent protocol, a peer-to-peer file-sharing protocol, to stream video in real time?[414] The answer depended on how network providers decided which applications were sufficiently alike to receive the same treatment. On the one hand, network providers could decide which applications are alike by focusing on the protocols used by the application and treat applications that use peer-to-peer file-sharing protocols differently from applications using other protocols.[415] In this case, they would treat Vuze like the other peer-to-peer file-sharing applications and slow it down. Alternatively, they could classify applications based on their sensitivity to delay. In this case, Vuze would be treated like other applications that stream video in real time and would not be slowed down. Like all applications that stream video in real time, Vuze is sensitive to delay. Thus, under the first approach, Vuze would perform worse during times of congestion than other

[410] Shaw (2006).

[411] See the discussion in footnotes 363 to 364and 381 above.

[412] For an overview of the practices, see Parsons (2009).

[413] See, e.g., Bell Aliant Regional Communications (2009), pp. 41-42, paras 85-87.

[414] Technically, YouTube (Ozer (2011)) and Vuze (P2PNet (2009)), like many other online video applications, use a technique called progressive download to create a near real-time streaming experience. See generally Odlyzko (2008); Ozer (2011).

[415] It seems that at least one deep packet inspection offering uses this approach. Cisco's protocol classification mechanism classifies peer-to-peer file-sharing applications such as BitTorrent based on the protocols used. It does not distinguish among time-sensitive and non-time-sensitive applications using the protocols. Cisco Systems (2012), pp. 10-12.

video applications like YouTube that also stream video in real time, but do not use peer-to-peer file-sharing protocols, putting Vuze at a competitive disadvantage.

The record of the proceeding did not resolve the question.[416] The concern that time-sensitive applications that use peer-to-peer file-sharing protocols may be harmed by practices that de-prioritize peer-to-peer file-sharing applications during times of congestion to improve the performance of time-sensitive applications has come up in other contexts as well.[417] As this example shows, network providers may deliberately or inadvertently choose a definition of "like" that distorts competition among applications within a class.

A network provider could also define classes of applications in a way that distorts competition among classes of applications. Again, this may happen deliberately or inadvertently. For example, network providers usually like the idea of providing low-delay service to online gaming.[418] Some online games are sensitive to delay, and charging the gamers for low-delay service would allow network providers to capture some of the value that online gamers realize from gaming.[419,420] By contrast, network providers seem to be less interested in providing low-

[416] In the CRTC hearings that were part of that proceeding, network provider representatives, when asked whether their traffic management system distinguished between time-sensitive and non-time-sensitive peer-to-peer applications, some network provider representatives testified that they excluded Skype, which has a peer-to-peer architecture, from their traffic management measures, because Skype, like all Internet telephony applications, is sensitive to delay (see, e.g., Canadian Radio-Television and Telecommunications Commission (2009c), para 4482-4483 (Telus); Canadian Radio-Television and Telecommunications Commission (2009d), paras 5640, 5953-5960 (Shaw) ; Canadian Radio-Television and Telecommunications Commission (2009e), paras 6253-6256 (Bell Canada). Contrary to Vuze, however, Skype does not use a peer-to-peer file-sharing protocol, so the treatment of Skype does not allow any conclusions regarding the treatment of Vuze. The responses cited above seemed to imply that traffic management practices apply to all peer-to-peer file-sharing applications, without distinguishing peer-to-peer file-sharing applications that are time-sensitive and those that are not. See also Canadian Radio-Television and Telecommunications Commission (2009b), paras 3938-3959 (Primus, peer-to-peer file-sharing is not time-sensitive).

[417] For example, in 2009, Vuze's General Counsel Jay Monahan asked the FCC to investigate the impact of Cox' trial network management system on peer-to-peer traffic, expressing concern that it would hurt the performance of Vuze. Cox' trial system de-prioritized applications such as peer-to-peer protocols (which, Monahan assumed, would include Vuze) that Cox had classified as non-time-sensitive during times of congestion to improve the performance of applications such as streaming video that Cox had classified as time-sensitive. P2PNet (2009);Cox Communications Inc. (2009). See also Public Interest Advocacy Centre (PIAC), et al. (2009), pp. 45-46, paras 126-130, pp. 51-52, paras 156-159, pp. 53-54, para 166;Vaxination Informatique (2009), p. 14, para 95, p. 21, paras 141-142; Distributel Communications Limited (2009), paras 10-16, 40-46 (all asking the CRTC to review its decision regarding Bell Canada's throttling of its wholesale ADSL access service on the grounds that CRTC based important parts of that decision on the assumption that all peer-to-peer file-sharing applications are not time-sensitive) and Canadian Radio-Television and Telecommunications Commission (2009a), paras 8-10 (rejecting the applications for review and reiterating the claim that peer-to-peer file-sharing applications are not time-sensitive).

[418] "Would this proposal allow us to offer Quality of Service to online games?" is usually one of the first questions I am asked when I discuss proposals for nondiscrimination rules with employees of network providers.

[419] See, e.g., Sandvine (2005), p. 1 ("Once aware of the amounts and types of gaming traffic on their network, a properly equipped service provider can prioritize it to create a gamer-friendly network. […] [S]ervice providers could perhaps offer their subscribers an optimized gaming tier with guaranteed QoS, while the escalating popularity of gaming drives demand for exactly that type of service. In this way, a MSO [Multi-system operator] can profit from an otherwise peripheral market trend." ibid.)

[420] Prices for Internet access service are currently independent of the application for which the service is used. For example, they do not change depending on whether a user is sending e-mail, placing a call, watching video or playing an online game. As I explained in my book, if network providers charge a uniform transport price and consumers value different applications differently, network providers will not be able to extract the full consumer surplus

delay service to Internet telephony applications like Skype or Vonage, since this would make these applications more competitive with the network providers' own telephony offerings.[421] Thus, a network provider may decide to offer low-delay service only to online gaming, but not to Internet telephony, arguing that these are different classes of applications because "gaming" and "telephony" are different uses of the Internet.[422] Internet telephony providers would argue that the correct class is "applications that are sensitive to delay," but their view would not matter until they had brought a complaint and succeeded in convincing the regulatory agency.

Finally, it may not be obvious which class an application belongs to, which allows network providers to inadvertently or deliberately hurt specific applications. A network provider may fail to provide the needed type of service to a certain application in a class because it does not realize that the application belongs to this class. For example, the Canadian Internet Service Providers in the example above may have intended to protect all real-time applications, even if these applications use peer-to-peer file-sharing protocols,[423] but may not have realized that there are applications like Vuze that use a peer-to-peer file-sharing protocol but are sensitive to delay. Alternatively, a network provider may argue that an application does not belong to a certain class, even if the network provider knows better, which would deprive the application of the needed type of service until the application provider has filed a complaint with the regulatory agency and succeeded.

Thus, even under rules that require like-treatment, network providers have ample discretion when defining classes of applications and assigning applications to classes. This allows network providers to use the provision of Quality of Service or other forms of differential treatment to deliberately harm certain applications or classes of applications. A benevolent network provider may inadvertently make decisions that have the same effect.

As the above examples show, disputes over which classes of applications are alike, or whether a certain application belongs to a certain class, are likely to be frequent and difficult to resolve, creating high costs of regulation.

User Choice

Apart from removing the application-agnosticism of the network, "like treatment" also violates the principles of user choice and innovation without permission.

---

associated with each application through the transport price alone (van Schewick (2010d), pp. 273-275). Under these circumstances, offering and charging for low-delay service for online gaming would allow network providers to extract some of the consumer surplus associated with online gaming that cannot be captured through the uniform transport price alone. Of course, network providers could also use this strategy to extract some of the consumer surplus associated with Internet telephony, but the trade-off would be more complicated: They would gain revenue from offering low-delay service to Internet telephony, but would lose revenue from their own telephony offering, because low-delay service increases the quality of Internet telephony relative to the network providers' offering, which makes Internet telephony relatively more attractive.

[421] This observation is based on conversations with network provider employees.

[422] For a real-world example of such an offering, see Sandvine (2013) (describing an Internet access plan by a South African DSL provider that prioritizes all ports used for online gaming, but not other ports or applications).

[423] See the responses of network provider representatives discussed in footnote 416 above.

Under "like treatment," network providers, not users choose which application should get which Quality of Service or differential treatment, violating the principle of user choice.[424] As I have explained elsewhere, the incentives of network providers and users are not necessarily aligned.[425] Network providers' incentive to reduce the performance of applications that may reduce their revenue from applications that are offered and provided separately from Internet access is one example of this phenomenon.[426] Thus, network providers do not always want to meet users' preferences. But even when they do, they may not be able to do so.[427] For example, if a network provider decides whether and when to offer QoS, it is forced to guess what the average user's priorities may look like, but these priorities may differ among users, and, for the same user, over time. In particular, a specific user's needs with respect to a particular application are not necessarily fixed:[428] A user's desire for QoS may differ considerably depending on the circumstances. For example, I may not care as much about the quality of my VoIP call when I am chatting with a friend as when I am doing a job interview. If I am playing a quick game at night I may be willing to tolerate a level of latency that I would not be willing to tolerate during an online gaming tournament. Normally, I may want file uploads to happen in the background and may want them to yield to other applications that are more important to me right now. But if I am uploading a large paper to a conference website just before the submission deadline, finishing this upload as quickly as possible will have the highest priority.[429] Thus, any QoS system that lets network providers determine whether and when to provide QoS may not be well aligned with users' needs. Network providers' attempts to determine which applications are time-sensitive and should receive special treatment during times of congestion will fail to meet users' needs for the same reasons.[430]

Innovation without Permission

Finally, "like treatment" harms application innovation by making it more difficult for new applications to get the type of service they need.[431] In order to get QoS, an application developer would have to convince network providers that its application belongs to a new class of applications that requires a certain type of service or that it is "like" an existing type of application that already receives that type of service, violating the principle of innovation without permission.[432] This introduces considerable transaction costs. Certain types of innovators (e.g. innovators that develop an application at home in their free time, non-commercial innovators, or

[424] van Schewick (2010f), p. 12.

[425] van Schewick (2010d), pp. 350-351.

[426] See footnotes 385 to 388 above and accompanying text.

[427] See also van Schewick (2010d), pp. 351.

[428] van Schewick (2008a), p. 7. See also Briscoe, Moncaster & Burness (2007), section 3.3; Riley & Scott (2009), p. 8; Lennett (2009), pp. 143-145; Free Press (2010), pp. 102-103; Yiakoumis, et al. (2012), Section 2.1.1. See also the discussion in footnote 142 above and accompanying text and in footnotes 481 to 482 and accompanying text below.

[429] This example is taken from Yiakoumis, et al. (2012), Section 2.1.1.

[430] See, e.g., Riley & Scott (2009), p. 8; Lennett (2009), pp. 143-144.

[431] van Schewick (2010f), pp. 12-13.

[432] Center for Democracy & Technology (2010), p. 29.

start-ups) may not have the resources necessary to engage in this type of negotiation with a potentially large number of network providers.[433] In addition, even if an innovator manages to contact a network provider, the innovator may not receive the appropriate QoS for its application if the innovator fails to convince the network provider. This is an example of the more general phenomenon that requiring cooperation or support from the network provider reduces the likelihood that innovative applications can be realized or successfully deployed.[434] Thus, requiring network providers to take action before an application can get the Quality of Service or differential treatment it needs violates the principle of innovation without permission and reduces the chance that new applications actually get the type of service they need.

Certainty and Costs of Regulation

In general, the rule is a lot clearer about which behavior is and is not allowed than the standards-based proposals discussed above. It clearly allows certain forms of Quality of Service, while banning others. In particular, the rule allows network providers to provide different types of service to different classes of applications that are not alike, as long as they do not discriminate among classes of applications that are alike or discriminate among like applications within a class. The rule does, however, prohibit network providers to offer a certain type of service only to some applications within a class. Thus, the rule restricts the evolution of the network more than approaches that allow all discrimination, but less than approaches that ban all discrimination.

With respect to specific instances of differential treatment among classes of applications, the rule provides less certainty than a more abstract reading of the provision may suggest. In particular, the ambiguities surrounding the definition of "like" make it difficult for network providers to predict whether their chosen definition will withstand regulatory scrutiny in case of a complaint. For the same reasons, application developers and their investors will not necessarily know in advance how far the rule's protections reach.[435] If adjudicators clarify the interpretation of "like" in the context of individual adjudications, this uncertainty may be reduced over time.[436] Until then, the rule will suffer from many of the problems associated with and will create similar social costs as the standards-based approaches discussed above, including high costs of regulation.

In sum, this rule is based on the assumption that discrimination among classes of applications that are not alike is socially harmless and should therefore be allowed. This

[433] Throughout the history of the Internet, many important innovations (including E-Bay, Facebook, Yahoo, Google, Apache Web Server, the World Wide Web, Flickr and Blogger) have been developed by innovators of this type. See van Schewick (2010h), pp. 3-5; van Schewick (2010d), pp. 204-213, 310-314, 318-328, 334-345 (discussing the importance of different types of low-cost innovators, including many examples).

[434] See footnote 51 above and accompanying text.

[435] On the importance of certainty for network providers and application developers, see footnotes 227 to 231 above and accompanying text.

[436] For a number of reasons, adjudicators may not necessarily have an incentive to clarify the meaning of key terminology beyond what is required to resolve the specific case under consideration. See the discussion in *Box 12: Will Individual Adjudications Reduce Uncertainty Over Time?* above.

assumption is not correct. In many cases, discrimination among classes of applications hurts some classes of applications, even if the classes are not alike. Like treatment removes the application-agnosticism of the network and violates the principles of user choice and innovation without permission. It allows network providers to deliberately or inadvertently distort competition among applications or classes of applications and interfere with user choice. Due to the ambiguities surrounding the definition of "like," the rule creates considerable uncertainty that will need to be resolved in case-by-case adjudications, resulting in social costs similar to the social costs of the standards-based approaches described above. Thus, like treatment creates considerable social costs and does not adequately protect the values that network neutrality rules are designed to protect.

Ban Application-Specific Discrimination, But Allow Application-Agnostic Discrimination

Instead, regulators or legislators should adopt a non-discrimination rule that clearly bans application-specific discrimination, but allows application-agnostic discrimination.[437,438] (Again, I use "applications" as shorthand for "applications, content, services, and uses.") Discrimination is application-specific if it is based on application or class of application, or, in other words, if it is based on criteria that depend on an application's characteristics ("application-specific criteria").[439] Application-specific criteria include what this paper calls "application" (i.e. the

[437] See van Schewick (2010e), pp. 6-8; van Schewick (2010f), pp. 13-16; van Schewick (2010b). In the Open Internet Proceeding, this proposal was supported by, e.g., networking experts (e.g., Reed (2010); NYSERNet (2010)); venture capitalists (e.g., Burnham (2010a); Wilson (2010)), entrepreneurs (e.g., Borthwick (2010); Srinivasan & Gupta (2010)) and non-profit organizations (e.g., The Council of Scientific Society Presidents (2010); North American Benthological Society (2010); Botanical Society of America (2010)).

[438] Many network neutrality proponents would couple a strict nondiscrimination rule that bans all discrimination with a reasonable network management exception that requires network management to be as application-agnostic as possible (e.g., Open Internet Coalition (2010), pp. 15-17, 49-50 (to qualify as reasonable network management, the practice must be narrowly tailored to address a legitimate network management purpose, i.e. in addition to other criteria, the practice must "result in as little discrimination or preference as reasonably possible" ibid., p. 49); Center for Media Justice, et al. (2010), p. 31-32, 35-41 (same). Thus, these proposals would allow application-agnostic differential treatment only as long as it was narrowly tailored to serve a legitimate network management purpose. The proposal described in the text goes beyond these proposals by allowing differential treatment based on application-agnostic criteria in general, not just when it is narrowly tailored to address a legitimate network management purpose.

[439] The two definitions of application-specific discrimination used in the text – "discrimination based on application or class of application" and "discrimination based on criteria that depend on an application's characteristics" – describe the same concept. In this paper, "application" refers to a specific instance of a specific type of application (see *Box 16: Terminology: "Application" and "Class of Application"* above). Thus, "discrimination based on application" is differential treatment of different instances of the same application type depending on which instance the user is using (e.g., Skype vs. Vonage). The specific instance of an application a user is using is also a characteristic of the application (i.e. it is a characteristic of the application whether it is Vonage or Skype).

A "class of applications" is a group of individual applications that share some common characteristic (see *Box 16: Terminology: "Application" and "Class of Application"* above). Thus, there are many different potential classes of applications based on which a network provider could discriminate, each defined by the criteria that are used to allocate the applications to the classes. For example, a class of applications may be the group of all applications of the same application type (e.g., Internet telephony, e-mail), all applications that use the same application-layer protocol (e.g., all applications that use SIP, all applications that use HTTP) or transport-layer protocol (e.g., all applications that use TCP, all applications that use UDP), or all applications that have similar technical requirements (e.g., all latency-sensitive applications, all latency-insensitive applications). A network provider discriminates "based

specific instance of an application a user is using, e.g., Vonage vs. Skype), application type (e.g. e-mail vs. Internet telephony), the application-layer protocol or transport-layer protocol the application is using (e.g. SIP vs. Skype's proprietary protocol, or TCP vs. UDP), or the application's technical requirements (e.g., latency-sensitive vs. non-latency-sensitive applications). (See also *Box 22: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination* below.) Since the term "applications" stands for applications, content, services or uses, the ban on application-specific discrimination applies equally to discrimination based on criteria that depend on characteristics of content or characteristics of a service or use. Thus, discrimination against certain content based on, e.g., publisher, author, content type, subject matter, or viewpoint would all be prohibited by the rule.

The rule should be coupled with an exception for reasonable network management which requires reasonable network management to be as application-agnostic as possible and allows the use of narrowly tailored application-specific measures only if a problem cannot be solved in an application-agnostic manner. (See *Box 21: The Exception for Reasonable Network Management* below.)

This rule plays an important role in the FCC's Open Internet order. The FCC's non-discrimination rule for fixed broadband access bans discrimination that is unreasonable. Whether discriminatory behavior complies with the rule just described (i.e. whether it is application-agnostic) is one of the factors the FCC will use to determine the reasonableness of discriminatory conduct under the non-discrimination rule and under the Open Internet Rules' exception for reasonable network management.

on class of application" if it treats the application differently depending on whether it belongs to the class or not. Since classes are defined by a common characteristic that the applications in the class share, discrimination based on class of application is the same as discrimination based on a characteristic of an application. For example, assume that a network provider discriminates against all applications that use the BitTorrent protocol. In this case, the criteria that is used to discriminate is "uses the BitTorrent protocol," which is a characteristic of an application. At the same time, the class is "all applications that use the BitTorrent protocol," and the network provider discriminates among applications based on whether they belong to this class or not.

Box 21
The Exception for Reasonable Network Management

Network neutrality rules usually include an exception for reasonable network management. Behavior that would otherwise violate the rule against blocking or the non-discrimination rule is allowed if it constitutes "reasonable network management" as defined by that exception.

The rule proposed in the text should be coupled with an exception for reasonable network management which requires reasonable network management to be as application-agnostic as possible and allows the use of narrowly tailored application-specific measures only if a problem cannot be solved in an application-agnostic manner.[440]

More formally, to qualify as reasonable network management, the practice would have to further a legitimate network management purpose and be narrowly tailored to address that purpose. In the context of network neutrality rules, the term "network management" refers to technical measures whose purpose is "to maintain, protect, and ensure the efficient operation of a network."[441] Network management includes, e.g., managing congestion or protecting the security of a network.[442] To qualify as narrowly tailored, the practice would have to, among other things, be as application-agnostic as possible and result in as little discrimination or preference as reasonably possible.[443] The treatment of network management practices under the proposed rule is described in more detail in the Section "Network Management" below.

The rule described in this section bans all discrimination among applications and classes of applications that is based on application-specific criteria, regardless of whether the applications or classes are alike or not. Thus, contrary to some non-discrimination rules in other areas of law, this approach does not require an analysis of whether the applications or classes of applications that are treated differently based on application-specific criteria are "alike" or "similarly situated." Nor is there an inquiry into whether the differential treatment of like applications or classes of applications is somehow justified. Instead, the rule strictly bans all discrimination based on application-specific criteria. The only way to justify instances of application-specific discrimination would be through the reasonable network management exception or any other exception that applies to the non-discrimination rule.

Under this approach, a network provider would not be allowed to treat Vonage differently from Skype, or Comcast's XfinityTV.com differently from Hulu. That would be discrimination

[440] See, e.g., van Schewick (2008a), pp. 4-8; van Schewick (2008b); van Schewick (2010k), p. 5; van Schewick (2010l); van Schewick (2010b). During the Open Internet Proceeding, the proposed exception was supported by, e.g., networking experts (e.g., NYSERNet (2010)); venture capitalists (e.g., Burnham (2010a); Wilson (2010)), entrepreneurs (e.g.,Srinivasan & Gupta (2010)) and non-profit organizations (e.g., The Council of Scientific Society Presidents (2010); North American Benthological Society (2010); Botanical Society of America (2010)).

[441] Center for Media Justice, et al. (2010), pp. 38-39.

[442] On the definition of network management, see also Center for Democracy & Technology (2010), pp. 41-43; (Center for Media Justice, et al. (2010), pp. 38-39).

[443] This formulation mirrors proposals by many network neutrality proponents. See, e.g., Open Internet Coalition (2010), pp. 48-50; Center for Media Justice, et al. (2010), pp. 35-41.

based on application.[444] Nor would it be allowed to treat online video differently from e-mail, treat applications that use the BitTorrent protocol differently from applications that do not use this protocol, or treat latency-sensitive applications differently from latency-insensitive applications. That would be discrimination based on class of application.[445] But it would be allowed to treat data packets differently based on application-agnostic criteria – criteria that have nothing to do with the application or class of application. (See also *Box 22: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination* below.) For example, a network provider could give one person a larger share of the available bandwidth if that person has paid for a higher tier of Internet service (e.g., if that person has paid for the "Up to 6 Mbps" Internet service packet instead of the "Up to 3 Mbps" Internet service packet).[446] During times of congestion, a network provider could give one person a larger share of the available bandwidth than another, for example, because this person pays more for Internet access or has used the Internet less over a certain period of time.[447] That would be application-agnostic discrimination. But it could not throttle the bandwidth available to a specific online video application such as Hulu in particular or to online video in general. That would be application-specific discrimination.

[444] On the meaning of "discrimination based on application," see footnote 439 above.

[445] On the meaning of "discrimination based on class of application," see footnote 439 above.

[446] See also Center for Democracy & Technology (2010), pp. 25-26.

[447] See also Center for Democracy & Technology (2010), pp. 25-26; Center for Democracy & Technology (2010), p. 19-20, 22 (arguing in favor of a strict non-discrimination rule, but asking the FCC to clarify that "the non-discrimination rule should not be interpreted to bar or restrict broadband providers from differentiating or prioritizing among Internet traffic based on the usage volumes, usage patterns, or subscription plans of the individual subscribers sending or receiving such traffic," ibid., p. 22).

Box 22

A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination

As David Reed has pointed out, thinking about the proposed rule in the context of the Internet's original architecture may help clarify the functioning of the rule. The Internet's original architecture was based on the layering principle and the broad version of the end-to-end arguments. The layering principle, as applied to networking, prescribes that a lower-layer protocol may not make any assumptions about the content or meaning of the message (or, more technically, protocol data unit) passed to it by a higher-layer protocol for delivery to its higher-layer protocol peer.[448] The lower-layer protocol may neither access nor act on the information contained in a higher-layer protocol data unit. This constraint preserves the central feature of layering: the independence of lower layers from higher layers. Thus, applied to the Internet Protocol, i.e., the protocol at the Internet layer, the layering principle prescribes that the Internet Protocol may not make any assumptions about the content or the meaning of the messages it is transporting on behalf of higher-layer protocols and may neither access nor act on the information contained in these messages. In other words, in the Internet's original architecture, the layering principle forced the network to be application-blind and application-agnostic.[449,450]

The proposed rule is not identical with the constraints imposed by the layering principle. In particular, the rule does not ban any violations of the layering principle as such. For example, the rule does not prevent network providers from using deep packet inspection in the network to access and analyze the higher-layer protocol data units transported by IP on behalf of higher-layer protocols; nor does it prohibit network providers from making assumptions about the meaning of the these higher-layer protocol data units. But if network providers have somehow acquired information about the content or meaning of the payload of the IP packets, the ban on application-specific discrimination prevents them from discriminating based on that information. Thus, with respect to network providers' ability to discriminate among applications based on application-specific criteria, the non-discrimination rule creates the same effect that compliance with the layering principle would

(continued)

[448] For a detailed explanation and references to the literature, see van Schewick (2010d), pp. 52, 56-57.

[449] See Reed (2010) (arguing that the Internet's original architecture forced the Internet to be application-agnostic).

[450] An application-blind network is unable to distinguish among the applications on the network, and, as a result, it is unable to make distinctions among data packets based on this information. An application-agnostic network may have information about the applications on the network, but does not make distinctions among data packets based on this information. Since the layering principle prevented the Internet layer from accessing higher-layer protocol data units or from making assumptions about their content and meaning, the Internet was unable to distinguish among the applications on the network and therefore, unable to make distinctions among applications. Thus, it was application-blind and, like all application-blind networks, necessarily application-agnostic. (See also *Box 4: Application-Agnostic v. Application-Blind* above.)

have created.[451] Translated into the less technical terminology used throughout this paper, while the rule creates the same effect as an application-blind network, the rule does not actually require the network to fully comply with the layering principle and be application-blind. It only requires the network to be application-agnostic.[452] (For an explanation of the rationale behind this decision, see *Box 4: Application-Agnostic v. Application-Blind* above.)

Thinking about the rule in this way may help clarify the distinction between application-specific and application-agnostic discrimination. In particular, in the current Internet, a practice that requires knowledge about the content and meaning of the payload of the IP packets traveling through the network would not be application-agnostic.[453] Note, though, that the classification of a practice as application-specific does not depend on how the knowledge about the content and meaning of the payload is acquired. For example, identifying encrypted Internet telephony applications or encrypted peer-to-peer file-sharing applications based on their traffic patterns and using that information as the basis for differential treatment constitutes application-specific discrimination. Similarly, identifying applications based on the port numbers typically used by that application and using that information as the basis for differential treatment constitutes application-specific discrimination as well.

Application-specific discrimination requires knowledge about the application or class of application that the user is using – knowledge that a network provider in an application-blind network (e.g., in the original Internet) would not have. A network provider in an application-blind network would, however, be able to engage in differential treatment that does not require knowledge about the application or class of application, and the rule maintains that ability. In sum, the rule bans all forms of discrimination that would not be possible in an application-blind network, and allows all forms of discrimination that would be available in such a network. Thus, the rule recreates through law the environment for application innovation and network use that an application-blind network such as the original Internet would create by virtue of its architectural design.[454]

[451] See also Reed (2010) (arguing that the rule proposed in the text effectively requires by law what the original design of the Internet required by code, and proposing to replace the proposed rule by one that "requir[es] those who offer Internet service to implement the Internet design as it was intended:" "In particular: We don't need a complex rule defining 'applications' in order to implement an application agnostic Internet. We have the basis of that rule – it's in the 'code' of the Internet. What we need from the 'law' is merely a rule that says a network operator is not supposed to make routing decisions, packet delivery decisions, etc. based on contents of the packet. Only the source and destination addresses and the labels on the packet put there to tell the network about special handling, priority, etc. need to be understood by the network transport, and that is how things should stay, if we believe that Barbara is correct that only application-agnostic discrimination makes sense." Ibid.)

[452] The decision to require only application-agnosticism in the context of network neutrality rules is not meant to imply that network providers should have the unlimited right to collect information about applications or user behavior. In particular, concerns about user privacy may justify limiting network providers' ability to collect information through privacy law. For further discussion, see *Box 4: Application-Agnostic v. Application-Blind* above.

[453] See Reed (2010).

[454] See also Reed (2010) (arguing that the rule proposed in the text effectively requires by law what the original design of the Internet required by code) and *Box 22: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination* above.

A network provider in an application-blind network cannot engage in application-specific discrimination because it does not have the information necessary to do so. By contrast, the non-discrimination rule proposed here does not prevent network providers from collecting application-specific information. It only prevents them from using this information to, e.g., discriminate among data packets or charge differently based on application-specific criteria.[455] Thus, while the rule creates the same environment for application innovation and network use as an application-blind network, it does not require the network to be "blind." It only requires the network to be application-agnostic.[456] (The rationale behind this decision is explained in *Box 4: Application-Agnostic v. Application-Blind* above. See also *Box 22: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination* above.)

Contrary to proposals based on an antitrust framework, the rule applies to all network providers, regardless of their market share in the market for Internet services,[457] and to all application-specific discriminatory conduct, regardless of whether the conduct is capable of monopolizing the market for the affected applications.[458] Any measure that singles out an application or class of applications for differential treatment tilts the playing field against some applications or classes of application and interferes with users' decisions about how to use the network, creating significant social costs.[459] The fact that the application-specific practice may serve a network provider's "legitimate business interest" as understood by the antitrust laws (e.g., if the goal of the practice is to manage congestion or to engage in price discrimination to recover the fixed costs of network infrastructure), is not sufficient to overcome the ban.[460] The social costs of application-specific discrimination result from the discriminatory conduct as such and are independent of the network provider's motivation. Even application-specific discrimination that does not seem to have the potential to harm any applications (e.g., providing QoS to different classes of applications according to their needs, or prioritizing time-sensitive applications over non-time-sensitive applications during times of congestion) creates considerable social costs.[461] At the same time, network providers can usually realize their legitimate goals using application-agnostic means that are not similarly harmful to application innovation, user choice, or the Internet's ability to reach its social, cultural or political potential. Network providers can, e.g., manage their networks in application-agnostic ways, price discriminate based on application-agnostic criteria or differentiate their services by offering

[455] For a full discussion of the kind of differential treatment subject to the rule, see the Section "Scope of Non-Discrimination Rules" above.

[456] The decision to require only application-agnosticism in the context of network neutrality rules is not meant to imply that network providers should have the unlimited right to collect information about applications or user behavior. In particular, concerns about user privacy may justify limiting network providers' ability to collect information through privacy law. For further discussion, see *Box 4: Application-Agnostic v. Application-Blind* above.

[457] See the discussion of the disclosure rule above. For a full analysis, see van Schewick (2010d), pp. 255-264. See also footnotes 197 to 198 and accompanying text.

[458] See footnotes 183 to 196 and accompanying text.

[459] van Schewick (2010d). See also footnotes (and accompanying text) 191 to 196, 368 to 378, 380 to 436 above.

[460] See footnotes 199 to 210 above and accompanying text.

[461] See footnotes 362 to 434 above and accompanying text.

Quality of Service in line with the rule.[462] In the rare cases in which a network management problem cannot be solved in an application-agnostic manner, the reasonable network management exception allows network providers to deviate from the non-discrimination rule in narrowly tailored ways.[463]

These rules are necessary because network providers' decisions about whether, when and how to engage in discrimination will not necessarily result in socially desired outcomes.[464] Network providers are not beneficial stewards of the Internet platform. They are private actors that pursue their private interests. Network providers' private interests often differ from users' interests, and even if they do not, network providers do not know what exactly users want.[465] Network providers' private interests and the public interests with respect to the evolution of the Internet diverge as well. It is this market failure that network neutrality rules are designed to address.[466] For a variety of reasons, network providers capture only a small part of the social value resulting from an open Internet. For example, they capture only some of the social benefits associated with application innovation or of the social benefits resulting from improved democratic discourse.[467] Moreover, most of the gains they are able to capture are uncertain and will be realized in the future, which leads network providers to discount them even more.[468] Thus, when network providers decide whether to discriminate among applications or classes of applications, the immediate private benefits of discriminating (i.e. the higher profits resulting from exclusionary conduct or from discriminatory network management) will often be higher than network providers' hyperbolically discounted share of the private benefits of refraining from discriminatory conduct.

Based on these insights, the rule takes away all the tools that would allow the network provider to deliberately or inadvertently interfere with competition and user choice – those involving application-specific discrimination –, while leaving the tools that cannot distort competition or violate user choice – those involving application-agnostic discrimination. By legitimizing a broad range of discriminatory conduct (that is, all conduct that is application-agnostic), the rule gives network providers flexibility to realize their legitimate goals such as network management, price discrimination or product differentiation, albeit through means that do not interfere with the values that network neutrality rules are designed to protect. If a network management problem cannot be solved in an application-agnostic manner, the reasonable

[462] On network management, see Section "Network Management" below. On Quality of Service, see Section "Quality of Service" below.

[463] See *Box 21: The Exception for Reasonable Network Management* above.

[464] This section summarizes arguments that I have developed in detail elsewhere. See the sources cited in the footnotes below.

[465] See footnotes 425 to 430 above and accompanying text.

[466] For a detailed discussion, see van Schewick (2010d), pp. 355-371 (describing the public interest), 371-375 (describing network providers' private interests and why they diverge from the public interest).

[467] van Schewick (2010d), pp. 373-374. See also Frischmann (2005); Frischmann & van Schewick (2007), pp. 400-403, 424-425; Hogendorn (2012).

[468] van Schewick (2010d), pp. 374-375.

network management exception provides a safety valve by allowing the network provider to deviate from the non-discrimination rule in narrowly tailored ways.

Protecting the factors that are at the core of the Internet's economic, social, cultural and political potential

Compared with other proposals for non-discrimination rules, the rule strikes the best balance between social benefits and social costs. The rule preserves the application-agnosticism of the network and the principle of user choice, two factors that have been central to the Internet's ability to foster innovation in the past.[469] By prohibiting application-specific discrimination, the proposed rule makes it impossible for network providers to distort competition among applications or classes of applications. The rule allows users, not network providers to choose how they want to use the network and which applications will be successful. Letting users make this choice not only increases the value of the Internet for users and for society, it is also an important part of the mechanism that enables application-level innovation to function effectively. In addition, maintaining application-agnosticism and user choice is crucial to allowing the Internet to realize its social, cultural and political potential.

Allowing the network to evolve

The proposed rule does not constrain the evolution of the network more than is necessary to reach the goals of network neutrality regulation. It provides room for networks to evolve.[470]

Quality of Service

The rule allows network providers to offer certain (though not all) forms of Quality of Service. In particular, it allows network providers to offer different classes of service if they meet the following conditions:

(1) the different classes of service are available equally to all applications and classes of applications;

(2) the user is able to choose whether, when and for which application to use which class of service;[471]

---

[469] On these factors and their economic, social, cultural and political impact, see footnotes 50 to 59 and accompanying text, *Box 4: Application-Agnostic v. Application-Blind* and *Box 3: The Importance of User Choice* above.

[470] For early versions of the arguments in this subsection, see van Schewick (2010e); van Schewick (2010f).

[471] Although the exact details vary (some would allow user-controlled Quality of Service only during times of congestion), many network neutrality proponents would allow the type of user-controlled Quality of Service described in the text. See van Schewick (2008a), pp. 7-8 (discussing the question in the context of the reasonable network management exception); Lennett (2009), pp. 143-145; Free Press (2010), pp. 103-104 (if there was a demonstrated need for priority, users should make the choice); Jordan & Ghosh (2010), pp. 12:14, 12:21-12:22; Center for Democracy & Technology (2010), pp. 26-27; Center for Democracy & Technology (2010), pp. 20-21, 23 (FCC should clarify that "the nondiscrimination rule shall not be interpreted to bar or restrict broadband providers from enabling individual subscribers to designate certain traffic streams for prioritized or differentiated treatment." Ibid., p. 23); Open Internet Coalition (2010), pp. 50-51; Skype Communications S.A.R.L. (2010), pp. 16-20; Google Inc. (2010), p. 36.

(3) the network provider is allowed to charge only its own Internet service customers for the use of the different classes of service.[472,473]

For example, a network provider could offer a low-delay service, a best-efforts service, a less-than-best-efforts service, and a guaranteed-bandwidth service. The decision of whether and when to use which service would be left to the user. For example, one user could use the low-delay service for Internet telephony, another may use it for online gaming, and a third user may use it for e-mail, if that is what that user wants. This type of user-controlled Quality of Service is technically feasible.[474,475] (The technical feasibility of this type of Quality of Service and other questions regarding the impact of the proposed rule on Quality of Service are discussed in *Box 23: Frequently Asked Questions on Quality of Service under the Proposed Non-Discrimination Rule* below.)

While the first two conditions directly flow from the proposed non-discrimination rule,[476] the third condition is based on additional considerations and would need to be encoded separately.[477]

A network provider who is allowed to charge for Quality of Service has an incentive to degrade the quality of the baseline, best-effort service to motivate users to pay for an enhanced type of service. The existence of this incentive is well-documented in the economic literature on price discrimination and one of the main motivations behind proposals to ban Quality of

---

See also IEEE-USA Committee on Communications Policy (2010), pp. 16-18, 23 (supporting the model of user-controlled QoS described in the text); IEEE-USA (2010), pp. 2-3.

[472] This condition would not constrain interconnection agreements in any way. Thus, payments among interconnecting networks would remain possible.

[473] I explain the rationale for this criterion in van Schewick (2010h); van Schewick (2010a), pp. 10-12. The question of whether and, if so, whom network providers should be allowed to charge for Quality of Service or other forms of preferential treatment is outside the scope of this paper. For a short overview of the options, see *Box 2: Charging for Quality of Service* and footnotes 23-24 above and accompanying text.

[474] This statement is based on many conversations with networking experts. See also Center for Democracy & Technology (2010), p. 27; Center for Democracy & Technology (2010), pp. 20-21; Jordan & Ghosh (2010), p. 12:21. Network providers sell business customers the option to choose the level of Quality of Service for their packets today. For example, customers of Verizon's Private IP Enhanced Traffic Management offering can mark their data packets for the desired class of service, which is then delivered by Verizon's network (Verizon (2006), pp. 25, 29; Verizon (2010), p.1). For a prototype targeting home users, see Yiakoumis, et al. (2012), Section 3, which won the Grand Prize at the Imagine App Challenge at the Cable Show 2012 (Silbey (2012); Spangler (2012)). For a more detailed discussion, see *Box 23: Frequently Asked Questions on Quality of Service under the Proposed Non-Discrimination Rule, Question 9* below.

[475] The form of user-controlled Quality of Service described in text does not violate the broad version of the end-to-end arguments. van Schewick (2010d), pp. 106-107. On the broad version of the end-to-end arguments, see footnote 2 above.

[476] Deviating from the first condition by making a specific type of service available only to some applications or classes of applications (e.g., only to the provider's own online video application, or only to online gaming, but not Internet telephony) would make distinctions among applications and classes of applications based on application-specific criteria (here: application or application type) and would thus violate the requirement that differential treatment must be application-agnostic. The second condition ensures that the differential treatment associated with the actual provision of the different types of services in the network happens based on an application-agnostic criterion (here: the type of service chosen by the user for that particular packet).

[477] See footnote 473 above.

Service.[478] To mitigate this problem, the rules should require the regulatory agency in charge of enforcing the network neutrality rules to monitor the quality of the baseline service and set minimum quality standards, if the quality of the baseline service drops below appropriate levels.[479]

This type of user-controlled Quality of Service offers the same potential social benefits as other, discriminatory or provider-controlled forms of Quality of Service without the social costs. In particular, it does not raise any of the problems associated with "like treatment." Contrary to like treatment, it preserves the application-agnosticism of the network, the principle of user choice, and the principle of innovation without permission:

First, the proposal maintains the application-agnosticism of the network: The provision of Quality of Service is not dependent on which applications users are using, but on the Quality-of-Service-related choices that users make. Thus, the network provider does not need to know anything about which applications are using its network in order for this scheme to work. The network provider only makes different classes of service available, but does not have any role in deciding which application gets which Quality of Service; this choice is for users to make. As a result, network providers cannot use the provision of Quality of Service as a mechanism to deliberately or inadvertently distort competition among applications or classes of applications.[480]

Second, since users choose when and for which applications to use a given type of service (in line with the principle of user choice), they can get exactly the Quality of Service that meets their needs.[481] As discussed repeatedly throughout this paper,[482] users' preferences with respect to Quality of Service will often differ across users and, for the same user, over time. Network providers may not always want to meet users' preferences and even if they do, they lack the information necessary to infer a specific user's preferences, which are often highly context-specific. As a result, Quality of Service systems that let network providers determine whether and when to provide Quality of Service will often fail to meet user needs. Letting users make the choice removes this problem.

[478] See footnotes 104 and 105 above and accompanying text.

[479] How to best address this incentive is an important question that deserves further discussion. The incentive to degrade the quality of the baseline service arises only if network providers are allowed to charge for Quality of Service. If they are not allowed to charge for it, they do not benefit from users' increased use of better-than-best-effort services and, therefore, do not have an incentive to degrade the quality of the baseline, best-effort service to motivate users to use more enhanced services. Thus, instead of adopting the solution proposed in the text, regulators could mitigate this problem by prohibiting network providers from charging for the provision of Quality of Service. Such a ban creates its own social costs, though. Like all issues related to charging for Quality of Service, these questions are outside the scope of this paper. For a discussion of the problem and of potential solutions, see van Schewick (2010a), pp.10-11. The European Union has adopted a similar rule following its review of the regulatory framework for telecommunications services. See Article 22(3) of the Universal Service Directive; European Commission (2007), pp. 92, 95-97, 101.

[480] See also Center for Democracy & Technology (2010), p. 27.

[481] van Schewick (2008a), p. 7 (discussing the question in the context of the reasonable network management exception); Lennett (2009), pp. 143-145; Riley & Scott (2009), p. 8.

[482] See footnotes (and accompanying text) 142, 424 to 430 above.

Third, in line with the principle of "innovation without permission," an innovator does not need support from the network provider in order for his application to get the Quality of Service it needs. The only actors who need to be convinced that the application needs Quality of Service are the innovator, who needs to communicate this to the user, and the user, who wants to use the application.[483] This greatly increases the chance that an application can get the type of service it needs.

User-controlled Quality of Service is not without costs. In particular, asking users to decide whether, when and for which applications to use Quality of Service imposes a burden on them. Users do not necessarily know enough to decide which class of service would be suitable for which application, so they may not be able to make "good" choices. Others may not want to bother with such technical details.

These problems can be mitigated, though. Applications know which type of service they would benefit from and could communicate this to the user. Applications that really benefit from a special type of service have an incentive to make it as easy as possible for their users to request that service. For example, if the network supports service discovery, the application could check whether the network offers potentially useful classes of service, monitor the performance of the application and ask for the user's permission to request an appropriate class of service, if the need arose. User interfaces can be designed to mask the technical details and make the selection of Quality of Service easy and intuitive.[484] For example, Skype could offer a simple button through which users could request high quality. The user only needs to understand that the button exists and that clicking on it results in a high-quality call. What the application does to make this happen (e.g., which type of service Skype requests from the network) can be hidden behind the interface. In addition to applications, other entities like end hosts or home routers could offer their own user interfaces that allow users to make Quality-of-Service-related choices and may even offer different interfaces to support users with varying levels of sophistication.[485]

Users who are not interested in making their own Quality-of-Service choices could outsource this task to a trusted third party. Home networks, which often support multiple devices, which run applications with diverse needs and are potentially operated by different users, have become quite complex, making them more difficult to manage and secure. Today, companies like Meraki, Powercloud or Aruba offer remote network management for small enterprises, and residential users may similarly benefit from the option to outsource the management of their home network to an outside provider.[486] Managing the user's Quality-of-

[483] See also Center for Democracy & Technology (2010), pp. 26-27.

[484] For a more detailed discussion of the role of user agents who bridge the gap between users and the network and a description of three prototypes for user agents with easy and intuitive user interfaces, see Yiakoumis, et al. (2012).

[485] See also *Box 23: Frequently Asked Questions on Quality of Service under the Proposed Non-Discrimination Rule, Question 9* below.

[486] See, e.g., Feamster (2010) (proposing an approach for outsourcing home network security), Yiakoumis, et al. (2011) (proposing a mechanism for outsourcing home network management more generally and describing an initial prototype that has been deployed).

Service needs by dynamically selecting appropriate classes of service for the various devices, applications and members of the user's household based on occasional or more frequent high-level input from the user might become part of such offerings.[487]

Thus, the increase in complexity can be mitigated, and any remaining costs are more than offset by the benefits that accrue to users and to society as a whole.

Box 23
Frequently Asked Questions on Quality of Service under the Proposed Non-Discrimination Rule

1. Does your rule require Internet service providers to offer Quality of Service?

The proposed rule does not require Internet service providers to offer Quality of Service; it only gives them the option to do so. But if a network provider wants to offer Quality of Service, it needs to offer it in compliance with the conditions described in the text.

2. Does your rule require Internet service providers to offer Quality of Service end-to-end, i.e. between the original source and ultimate destination of data, across the networks of different providers ("end-to-end, inter-provider Quality of Service")?

The rule does not require Internet service providers to offer Quality of Service end-to-end.[488] Network neutrality rules usually have a limited scope, and the obligations imposed by the non-discrimination rule, including any constraints on the provision of Quality of Service, do not go beyond the scope of the rules. For example, the FCC's Open Internet Rules only apply to providers of broadband Internet access service as specified by the rules.[489] According to the text of the Open Internet Order, the Open Internet Rules do not apply to "Internet backbone services (if those services are separate from broadband Internet access service)"[490] and "[the] rules apply only as far as the limits of a broadband provider's control over the transmission of data to or from its broadband customers."[491]

(continued)

[487] See Yiakoumis, et al. (2011) (explicitly discussing this possibility).

[488] The rule does not require Internet service providers to offer Quality of Service at all, but if they offer it, they do not have to offer it end-to-end across multiple networks.

[489] 47 C.F.R. §8.11 and Federal Communications Commission (2010c), pp. 28-32, paras 44-52.

[490] Federal Communications Commission (2010c), p. 29, para 29.

[491] Federal Communications Commission (2010c), p. 29, footnote 150.

3. How does your rule constrain the provision of end-to-end, inter-provider Quality of Service?

Although network neutrality rules may not cover the entire path of traffic between two endpoints and therefore cannot influence network providers' actions along the entire path, network neutrality rules constrain the provision of end-to-end, inter-provider Quality of Service along the part of the path that is subject to the rules. Thus, an Internet service provider subject to network neutrality rules that implemented my proposal would have to comply with the restrictions on the provision of Quality of Service outlined in the text on the portion of the path that is subject to the rules, regardless of whether the provision of Quality of Service was restricted to the Internet service provider's network or whether the provision of Quality of Service was part of an end-to-end, inter-provider offering.

Although end-to-end, inter-provider Quality of Service has not been widely deployed in the public Internet for technical and commercial reasons,[492,493] any network neutrality non-discrimination rule should not prevent a migration to end-to-end, inter-provider Quality of Service in the future. My proposal reflects this goal. While the proposal constrains an Internet service provider's ability to charge the end users – the ultimate source and destination of the traffic for which Quality of Service is desired – for the provision of Quality of Service, it does not constrain payment arrangements for Quality of Service in the context of interconnection agreements. If a subscriber ("user A") requests a certain class of service for particular traffic (whether upstream or downstream), its Internet service provider ("ISP A") is allowed to charge only its own subscriber, i.e. user A, for the provision of that class of service on ISP A's network. It is not allowed to charge the end user at the other end of the connection (e.g., an application or content provider who is not a customer of ISP A's Internet access service) for the provision of Quality of Service on ISP A's network.[494] By contrast, if an Internet service provider and an interconnecting network wanted to exchange compensation as part of an interconnection agreement that involved the provision of

(continued)

[492] For an analysis of the reasons, see, e.g., Armitage (2003); Bell (2003); Burgstahler, et al. (2003); Huston (2012).

[493] Large carriers such as Verizon or Deutsche Telekom offer business customers private IP services that provide different classes of service end-to-end between a customer's corporate networks in different countries. In these cases, the carrier controls the provision of different classes of service either directly (if the customer's networks are attached to networks directly controlled by the carrier) or indirectly (if the customer's networks are attached to networks of a different carrier with which the first carrier (e.g., Verizon or Deutsche Telekom) has interconnection agreements that include service level agreements for different classes of service). See, e.g., Verizon (2006), pp. 41-44; Verizon (2010), p. 1; von Bornstaedt (2012).

[494] See footnote 473 above and accompanying text.

different classes of service, my proposal would not prevent them from doing so. The proposal explicitly states that the proposal's restrictions on charging end users for Quality of Service do not constrain interconnection agreements in any way.[495]

4. If the proposed rule does not require Internet service providers to offer end-to-end, inter-provider Quality of Service and end-to-end, inter-provider Quality of Service is not currently widely deployed in the public Internet, what, if any, is the benefit of Quality of Service offerings that are limited to an Internet service provider's network?

Naturally, a Quality of Service offering that is limited to an Internet service provider's network cannot provide the desired class of service for the whole path of traffic between a subscriber's end host and the other end host involved in the connection if the other end host connects to the Internet through another Internet service provider.

A limited offering may nevertheless provide benefits in two cases: First, a limited offering is equivalent to an end-to-end offering if the traffic does not leave the Internet service provider's network. Second, providing Quality of Service only on the Internet service provider's network can be beneficial if the main points of congestion are on that network. In the US and Europe, the access networks are the main sources of congestion, while the backbone is not congested.[496,497] Thus, two users who are talking to each other via an Internet telephony application and subscribe to different Internet service providers may encounter congestion on both access networks, but not on the backbone. Assume that each of the two Internet service providers is offering low-delay service for upstream traffic (i.e. traffic from the user to the Internet) and for downstream traffic (from the Internet to the user) between the user's end host and the edge of the Internet service provider's access network.[498] Under these conditions, each user could protect the call from the impact of potential congestion on the network of its own Internet service provider by choosing low-delay service for the corresponding traffic: User A's choice of low-delay service would protect the upstream and downstream portions of the call on user A's access network, while user B's choice of low-delay service would protect the upstream and downstream portion of the call on user B's access network. If there is no congestion on the backbone network, the lack of low-delay service on the backbone would not affect the quality of the call in any way.

(continued)

[495] See footnote 472 above.

[496] Bauer, Clark & Lehr (2009), p. 16 (arguing that "we expect that, at least for the near term, the access networks will remain the dominant constraint on achievable throughput" given the "relative economics" of adding capacity in the backbone versus in the access networks).

[497] There may be other points of congestion. For example, the links that connect interconnecting networks are often congested, too.

[498] On providing Quality of Service for upstream and downstream traffic, see Question 9 below.

Thus, if the main points of congestion are on the access networks, Quality of Service offerings that are limited to the access networks will improve the quality of applications that benefit from the offered classes of service even if these classes of service are not offered end-to-end.

Providing Quality of Service over a limited domain only is also in line with standards of the Internet Engineering Task Force. DiffServ, one of two architectures for the provision of Quality of Service standardized by the Internet Engineering Task Force, was explicitly designed to allow the provision of Quality of Service within a particular network or set of networks only.[499] Adjacent DiffServ-enabled networks can then be combined to provide Quality of Service across larger parts of the Internet until so many networks are DiffServ-enabled that it is possible to provide Quality of Service end-to-end.[500]

5. Under your rule, what prevents an individual user from marking all of his packets as low-delay traffic?

Nothing in the rule requires network providers to allow their Internet service customers to use an unlimited amount of a specific class of service. Under the proposed rule, network providers can impose limits on the use of a specific class of service, as long as the limit is application-agnostic. For example, they might sell subscribers the right to use up to x MB of low-delay service over a certain period of time (e.g., per second).[501] When a subscriber's traffic that is marked for low-delay service enters the network provider's network, the network checks whether the amount exceeds the contractually specified limit on low-delay traffic for that subscriber and remarks any excess packets according to an agreed upon scheme. For example, the provider may remark packets that exceed the maximum limit on low-delay traffic as best-effort service.[502] Under such a scheme, an individual subscriber would not be able to send more than the specified maximum amount of low-delay traffic into the network. Imposing a maximum limit on the amount of traffic of a specific type of service without tying it to the use of a specific application or class of application is application-agnostic, and thus compliant with the proposed rule.

(continued)

[499] Carpenter & Nichols (2002), p. 1481; Blake, et al. (1998), Section 1.3, p. 8.

[500] Carpenter & Nichols (2002), p. 1484, 1487-1488.

[501] In real life, the contract may specify traffic profiles in more detail. For example, a contract that allows a user to mark packets for a specific class of service may specify the maximum average rate, peak rate and burst size at which traffic marked for that class of service is allowed to enter the network. See Kurose & Ross (2010), pp. 657-659.

[502] For example, Verizon's private IP offerings for companies allow Verizon's customers to mark each packet with the class of service desired for that packet. When the traffic enters Verizon's network, the network polices the traffic to ensure it conforms to the contractually specified traffic profile for the relevant class of service. See Verizon (2006), p. 29. On policing and Quality of Service more generally, see Kurose & Ross (2010), pp. 650-652, 657-659, 663; Peterson & Davie (2012), p. 550.

By contrast, network providers are not allowed to specify maximum limits for a certain class of service that differ depending on which application is using the class of service (e.g., the subscriber is allowed to use up to x MB of low-delay service per second for Internet telephony, but only y MB of low-delay service per second for online gaming). Such application-specific limits on the use of a specific class of service would violate the proposed rule.

6. What happens if all users want to use their contractually specified maximum amount of low-delay service at the same time?

Today's access networks are built on the assumption that not all subscribers use the contractually specified maximum amount of bandwidth at the same time.[503] It is possible that the provisioning of capacity for certain classes of service would share this characteristic. In such a network, the total amount of traffic desiring, e.g., low-delay service may exceed the capacity available for that service if too many subscribers simultaneously send the maximum amount of traffic marked for low-delay service into the network. In this situation, the rule allows network providers to allocate the available capacity for low-delay traffic using application-agnostic criteria. For example, the network could give a relatively higher share of the available capacity to users who have used the low-delay service less over a certain period of time, or the provider could sell users the right to use relatively more low-delay capacity during times of congestion. These would be application-agnostic criteria.

By contrast, network providers would not be allowed to give a smaller share of the overall capacity for low-delay service to users who use the low-delay service for online video than to users who use the low-delay service for Internet telephony or online gaming. These would be application-specific criteria for allocating capacity during times of congestion, which would violate the proposed rule.

7. Would your rule allow admission control as part of the provision of Quality of Service?

Many Quality of Service architectures include admission control: When a new data flow requests a particular class of service, the network checks whether it has the resources necessary to provide that class of service. If it does have the necessary resources, it admits the flow. If does not have them, it rejects it. Admission control prevents situations where all flows are admitted to the network, but none receives the amount of resources necessary to perform adequately.[504]

(continued)

[503] This practice is called oversubscription. For a short explanation of the practice, see St. Arnaud (2009), paras 8-15.

[504] On admission control in Quality of Service architectures, see Peterson & Davie (2012), pp. 539, 542-543 and Kurose & Ross (2010), pp. 665-669.

Admission control would not violate the proposed rule as long as the decision to admit a new flow is based on application-agnostic criteria. By contrast, a network could not use application-specific criteria. For example, it could not base admission of competing flows on the applications that the flows belong to.

8. Under your rule, can the provision of Quality of Service ever be based on application-specific criteria?

The proposed non-discrimination rule is subject to a reasonable network management exception and to any other exceptions specified in the actual rules. Thus, an Internet service provider could use application-specific criteria for the provision of Quality of Service or for admission control if the conditions underlying one of these exceptions are met. For example, a network provider could give network management traffic and routing traffic precedence over other traffic if the conditions of the reasonable network management exception are met. Similarly, the FCC's Open Internet Rules allow Internet service providers to prioritize emergency communications.[505]

9. Is the type of user-controlled Quality of Service you describe in the text technically feasible?

This type of user-controlled Quality of Service is technically feasible.[506] Network providers sell business customers the option to choose the level of Quality of Service for their packets today. For example, customers of Verizon's Private IP Enhanced Traffic Management offering can mark their data packets with the desired class of service, which is then delivered by Verizon's network.[507]

If an Internet service provider offers Quality of Service in compliance with the proposed rule, the network offers different classes of service, while the user decides whether, when and for which application it would like to use a specific class of service. The network then treats the user's traffic according to the user's choices within the constraints (e.g., contractually specified limits on the use of specific types of service or admission control) described in Questions 5-7 above.

(continued)

[505] Federal Communications Commission (2010c), pp. 59-60, paras 108-110 and 47 C.F.R. § 8.9 ("Nothing in this part supersedes any obligation or authorization a provider of broadband Internet access service may have to address the needs of emergency communications or law enforcement, public safety, or national security authorities, consistent with or as permitted by applicable law, or limits the provider's ability to do so.")

[506] This statement is based on many conversations with networking experts. See also Center for Democracy & Technology (2010), p. 27; Center for Democracy & Technology (2010), pp. 20-21; Jordan & Ghosh (2010), p. 12:21.

[507] Verizon (2006), pp. 25, 29, 56. For a prototype targeting home users, see Yiakoumis, et al. (2012), Section 3, which won the Grand Prize at the Imagine App Challenge at the Cable Show 2012 (Silbey (2012); Spangler (2012)).

For this to work, (1) the user must be able to express its Quality-of-Service-related choices through an appropriate user interface, and (2) these choices must be signaled to the network so that the network can treat the affected traffic accordingly. To provide maximum flexibility for innovation, the proposed rule deliberately does not prescribe how the user interface or the signaling should be implemented.

Technically, the user interface and the signaling could be implemented in a number of ways: For example, Internet service providers (in their offerings for business customers) or corporate intranets often use DiffServ to provide different classes of service. Under the DiffServ standard, each IP data packet carries information that indicates the class of service requested by that packet.[508] On the user's side, different entities – e.g., applications, the user's end host or a home router – could expose a user interface that allows the user to select the desired class of service for specific applications or data flows.[509]

For upstream traffic (i.e. traffic from the user's end host into the Internet), the entity that offers a user interface for selecting the classes of service – e.g., the application, the user's end host or a home router – could directly set the Quality-of-Service-related bits in affected packets in line with the user's choice as expressed through that interface.[510]

For downstream traffic (i.e. traffic from the Internet to the user), the user (or, more precisely, the entity responsible for signaling the user's Quality-of-Service-related choices to the network) cannot directly set the Quality-of-Service-related bits in affected packets since it does not control the router at which the packets enter the

(continued)

[508] The information about the requested class of service is called Differentiated Services Code Point (DSCP) and is encoded in the Differentiated Services Field. See Blake, et al. (1998) and Nichols, et al. (1998).

[509] As explained in the main text of the paper, such entities could offer different user interfaces for users with varying levels of sophistication. Interfaces could be designed to make the choice intuitive and easy for average users or offer more granular control for users with more expertise. For an exploration of the design space for user agents (i.e. entities that translate the user's intent into technical requests to the network), see Yiakoumis, et al. (2012).

[510] Applications, end host or home router could all expose an interface for selecting classes of service, and each could set the appropriate Quality-of-Service-related bits in the affected upstream packets (i.e. packets from the end host to the Internet). While the different options are likely to have different advantages and disadvantages, a discussion is outside the scope of this paper.

domain over which the user's Internet service provider offers Quality of Service.[511] Thus, there needs to be a mechanism that allows the user or the entity acting on the user's behalf to signal the class of service desired for particular inbound traffic (e.g., for the traffic of an Internet telephony connection flowing from the other party towards the user or for online gaming traffic traveling from the game server to the user) to the device at which the traffic enters the domain over which the user's Internet service provider offers Quality of Service. Again, there are various options for implementing that signaling. For example, the Next Steps in Signaling Framework standardized by the Internet Engineering Task Force could be used to solve this problem.[512] The Center for Democracy and Technology described another potential mechanism in its reply comments in the FCC's Open Internet proceeding.[513] Using yet another mechanism, a recent prototype sends a separate signaling message directly to the Internet service provider's network. The signaling message specifies the desired class of service for a specific upstream and/or downstream data flow or set of data flows.[514]

In the examples discussed so far, the user makes its choices through user interfaces exposed by applications, end hosts or home routers which are then signaled to the network. Alternatively, an Internet service provider could offer its own user

(continued)

[511] Data packets traveling from the Internet towards the user (receiver) for which the user desires a certain class of service are unlikely to carry the correct information about the desired class of service when they enter the domain over which the user's Internet service providers offers Quality of Service, even if the sender has marked the packets with the class of service desired by the receiver. To see this, assume that the sender has marked the packet with the class of service desired by the receiver. In today's Internet, a data packet usually traverses a number of networks on its path from its original source to its ultimate destination. If that marking was preserved as the packet traveled across the Internet, it would still be there when the packet entered the receiver's Internet access network. In today's Internet, however, when a packet passes from one network A to the next network B, network B usually overwrites any DSCP marking that might have been present, unless the two networks have an interconnection agreement that includes the provision of different classes of service. Thus, even if the sender marked the packet with the desired class of service when it sent off the packet, the packet is unlikely to carry the correct information about the desired class of service when – after having traversed other networks – it finally enters the domain over which the receiver's Internet service provider offers Quality of Service.

[512] See Hancock, et al. (2005); Manner, Karagiannis & McDonald (2010).

[513] Center for Democracy & Technology (2010), p.21: "CDT believes that engineers could devise solutions to allow user-directed prioritization of downstream traffic as well. For example, upstream traffic that a user sends to a particular online service could be marked for priority with an encrypted token generated by the broadband provider; the online service, in sending its response, could copy that encrypted token to mark the downstream traffic for priority as well. The broadband provider would recognize the encrypted token as an authentic indication of a user's prioritization request. Other approaches could be possible as well; the point is that providing effective user-directed priority should not pose any insurmountable technical challenge."

[514] Yiakoumis, et al. (2012), Section 3.1 ("User-agents communicate with the infrastructure using an out-of-band messaging scheme, through which they define a set of flow to service mapping. On the infrastructure side, the messages are received by an OpenFlow controller which enforces the mappings to appropriate queues.") and Yiakoumis (2013) ("I signal the ISP using a separate message, like the following: "all traffic to/from IP address 10.10.20.54, port 80 and me requires type of service X." This is directly mapped to a simple configuration to the ISPs router(s).").

interface (e.g., in form of a web site) that allows the user to choose the desired class of service for its various applications and then treat the user's traffic according to these choices.

10. I like your non-discrimination rule, but I am not convinced that network neutrality rules should allow Quality of Service. Can I adopt your non-discrimination rule, but ban Quality of Service? And if I do so, would your non-discrimination rule still be useful?

Non-discrimination rules apply to any form of differential treatment, not just to the differential conduct necessary to provide Quality of Service.[515] For example, non-discrimination rules govern any differential handling of packets within the network including, e.g., the allocation of bandwidth among users during times of congestion or other differential treatment of data packets to manage congestion. Non-discrimination rules also affect, e.g., whether network providers can count only some traffic, but not other towards users' monthly usage caps, what factors network providers can use to price discriminate, or whether they can charge different Internet access prices depending on the application used, independent of the traffic created by the application. Thus, the choice of non-discrimination rule has important implications for many questions other than whether and under which conditions to allow Quality of Service.

The non-discrimination rule proposed in the text allows certain forms of user-controlled Quality of Service. Regulators who like the rule, but disagree with the paper's arguments regarding Quality of Service, could adopt the non-discrimination rule proposed here and couple it with an explicit ban on Quality of Service. In this case, the non-discrimination rule would apply to any form of differential treatment except for Quality of Service. Thus, the choice of non-discrimination rule can be separated from the decision whether to allow Quality of Service if that is desired.

Network Management

The proposed non-discrimination rule allows any differential treatment that is application-agnostic. This includes any network management practices that treat traffic differently based on application-agnostic criteria.[516] In addition, the proposed reasonable network management

[515] For a more detailed description of the scope of non-discrimination rules, see Section "Scope of Non-Discrimination Rules" above.

[516] Network management practices that treat traffic differently based on application-agnostic criteria would be allowed under the proposed non-discrimination rule as such, since that rule allows any differential treatment that is application-agnostic. Thus, under the proposed rule, such application-agnostic network management practices would not need to meet the requirements of the reasonable network management exception. See footnote 438 above. On the definition of network management, see *Box 21: The Exception for Reasonable Network Management* above. For a longer discussion of the policy arguments driving the treatment of network management measures proposed in the

exception requires network management to be as application-agnostic as possible. The exception allows the use of narrowly tailored application-specific measures only if a problem cannot be solved in an application-agnostic manner.[517]

This proposal gives network providers the tools they need to manage their networks and maintain the quality of the Internet experience for all users, while preserving the application-agnosticism of the network and user choice as much as possible.[518] Network providers can enforce fairness among users and prevent aggressive users from overwhelming the network by allocating bandwidth among users in application-agnostic ways. During times of congestion (i.e., during times when a link's average utilization is high),[519] network providers may limit the amount of capacity available to users of that link based on application-agnostic criteria. For example, a network provider could give one person a larger share of the available bandwidth than another, for example, because this person pays more for Internet access or has used the Internet less over a certain period of time.[520] Even during times of congestion, network providers cannot, however, interfere with how users use the (limited) capacity available to them, e.g., by selectively blocking or discriminating against specific applications or classes of applications. Such application-specific traffic management practices would violate the rule's ban on application-specific discrimination and the reasonable network management exception. Thus, while the amount of bandwidth available to a user during times of congestion may be limited, users still get to decide how to use that bandwidth without interference from network providers.

To the extent that applications benefit from relative prioritization or other forms of differentiated treatment during times of congestion (i.e., during times when a link's average utilization is high),[521] network providers could allow users to choose which applications to prioritize or otherwise treat differently during these times.[522] As long as the option to be prioritized or be treated differently is offered equally to all applications or classes of applications (i.e. not tied or restricted to specific applications or classes of applications) and the choice of which applications to prioritize or treat differently is left to the user, this form of network

---

text, see van Schewick (2008a), pp. 4-8; van Schewick (2008b); van Schewick (2010b). See also footnotes (and accompanying text) 207 to 210, 245 to 252 and 262 to 269.

[517] For a more detailed description of the reasonable network management exception proposed by this paper, see *Box 21: The Exception for Reasonable Network Management* above.

[518] Application-agnostic network management preserves application-agnosticism and user choice. These principles are compromised only if application-agnostic network management is impossible.

[519] In discussions of the reasonable network management exception, the term "congestion" is generally used according to the definition of congestion used by network providers. Under that definition, congestion occurs if the average utilization of a link over a certain time period exceeds a certain threshold. See *Box 8: Definitions of Congestion and Benefits from Quality of Service* above.

[520] See also Center for Democracy & Technology (2010), pp. 25-26; Center for Democracy & Technology (2010), p. 19-20, 22 (arguing in favor of a strict non-discrimination rule, but asking the FCC to clarify that "the non-discrimination rule should not be interpreted to bar or restrict broadband providers from differentiating or prioritizing among Internet traffic based on the usage volumes, usage patterns, or subscription plans of the individual subscribers sending or receiving such traffic," ibid., p. 22).

[521] On the use of the term congestion, see footnote 519 above.

[522] van Schewick (2008a), pp. 7-8; Jordan & Ghosh (2010), pp. 12:17-12:20.

management would be consistent with the non-discrimination rule and reasonable network management exception proposed above.[523]

Tools for application-agnostic congestion management are available today. For example, Comcast, the largest provider of broadband Internet access services in the US,[524] adopted an application-agnostic congestion management system in response to the FCC's Order against Comcast in 2008.[525] According to Comcast, "Comcast's trials and subsequent national deployment indicate that this new congestion management system ensures a quality online experience for all of Comcast's HIS [High Speed Internet] customers."[526] Thus, it is possible to protect the quality of the Internet experience of all Internet service customers in application-agnostic ways. Beyond Comcast's approach, vendors have developed network management solutions that allow the network provider to allocate bandwidth among users in an application-agnostic manner, while letting users choose the relative priority of applications within the bandwidth allocated to them.

The proposed rule is also compatible with new standards that are currently being developed by the Congestion Exposure Working Group in the Internet Engineering Task Force.[527] These standards would evolve the existing standards for the TCP/IP protocol suite in a way that allows the network provider to determine how much a user's traffic is contributing to congestion at any point in time. This information would allow network providers to manage their networks based on a user's contribution to congestion.[528] Network providers could use this information, e.g., to allocate bandwidth among users during times of congestion based on their contribution to congestion, charge users based on their contribution to congestion, or to count only traffic that contributes to congestion towards a user's monthly usage cap. Since a user's contribution to congestion is an application-agnostic criterion, all of these forms of differential treatment would be allowed under the proposed rule.

From a technical perspective, application-agnostic network management has the added advantage of ending the arms-race between application developers, users and network providers that often develops in networks that use application-specific network management practices. Network management practices that single out specific applications or classes of applications for special treatment often motivate application developers to masquerade their applications to evade performance-reducing practices targeting their applications or to take advantage of performance-enhancing treatment provided to other applications, resulting in a cat-and-mouse game between network providers on the one hand and application developers

[523] See the discussion of Quality of Service in the previous section.

[524] Leichtman Research Group (2014).

[525] For descriptions of Comcast's application-agnostic network management system, see Comcast Corporation (2008b); Zachem (2009); Bastian, et al. (2010).

[526] Bastian, et al. (2010), Section 8.

[527] See Internet Engineering Task Force (2010).

[528] For an overview, see Briscoe, Woundy & Cooper (2012); Jacquet, Briscoe & Moncaster (2008).

and users on the other hand.[529] Application-agnostic network management practices remove this incentive, freeing resources at network providers, application developers and users.

In sum, network providers will often be able to manage their networks in application-agnostic ways, which maintains the application-agnosticism of the network and the principle of user choice. In the rare cases in which a problem cannot be solved in an application-agnostic manner, the reasonable network management exception provides a safety valve by allowing network providers to use narrowly tailored application-specific measures.[530]

Certainty and Costs of Regulation

The proposed rule does not suffer from the same definitional ambiguities and does not offer similar possibilities to game the system as a rule that requires like treatment, resulting in lower costs of regulation. Since the rule clearly specifies in advance which behavior is and is not acceptable, it is also easier and less expensive to enforce than the standards-based approaches discussed above.[531] Contrary to those approaches, the rule is immediately applicable to all industry participants. This not only removes the need to re-adjudicate similar cases again and again.[532] It also avoids the intertemporal inconsistencies across industry actors and across different decision-makers that are unavoidable under a standards-based approach.[533]

Of the approaches discussed in this paper, only the all-or-nothing approaches and the disclosure rule have lower costs of regulation. They are, however, either too restrictive, prohibiting socially beneficial forms of discrimination and restricting the evolution of the network more than necessary to protect the values that network neutrality regulation is designed to protect (like the approaches that ban all discrimination) or do not sufficiently protect the values that network neutrality regulation is designed to protect (like the approaches that allow all discrimination or the disclosure rule).

By clearly specifying acceptable and unacceptable behavior in advance, the rule provides certainty to all industry participants and avoids the many problems associated with determining the legality of specific discriminatory conduct after the fact in case-by-case adjudications outlined above. In particular, it does not tilt the playing fields against those – end users, low-cost innovators or start-ups – who do not have the resources to fight over the correct interpretation and application of the rule in the future. Network providers know how they can manage their networks. Application developers and their investors know that they will have a fair

[529] See footnotes 263 to 266 above and accompanying text.

[530] For a more detailed description of the reasonable network management exception proposed by this paper, see *Box 21: The Exception for Reasonable Network Management* above.

[531] See generally Pierce (2010), pp. 497, 499-500. For a more detailed discussion, see the analysis in Sections "Ban Discrimination That is Unreasonable" and "Problems with Case-by-Case Adjudications" above.

[532] See generally Pierce (2010), pp. 497-500.

[533] See generally Pierce (2010), pp. 500-501.

chance in the market place: that they will be able to reach users and compete with other applications on the merits, without interference from network providers.[534]

In addition to the costs of the non-discrimination rule itself, costs will be created by the proposed exception for reasonable network management and the proposal to require the regulatory agency in charge of enforcing the network neutrality rules to monitor the quality of the baseline service and set minimum quality standards, if the quality of the baseline service drops below appropriate levels.

The reasonable network management exception may require case-by-case adjudications to determine whether application-agnostic ways of solving the network management problem in question are available and, if not, whether the chosen application-specific measure is narrowly tailored. All proposals for network neutrality rules include an exception for reasonable network management to ensure that network neutrality rules do not unduly interfere with network providers' ability to manage their networks. Thus, all network neutrality regimes will be afflicted with the costs of adjudications under that exception. Contrary to some alternative proposals for a reasonable network management exception which would allow network management as long as it is "reasonable" without further specifying the term,[535] the proposed exception clearly specifies the standard that will be used to determine whether a specific measure constitutes reasonable network management. Thus, the exception provides more certainty to industry participants and will be easier and less expensive to enforce than exceptions that leave both the development of the standard and the application of the standard to case-by-case adjudication.[536]

The proposal to task the regulatory agency with monitoring the quality of the baseline service and to set minimum quality standards, if necessary, is designed to address network providers' incentive to degrade the quality of the baseline, best-effort service in order to motivate users to pay for an enhanced type of service. This incentive exists in all network neutrality regimes that allow network providers to charge for the provision of Quality of Service,[537] so all network neutrality regimes that allow charging for Quality of Service will have to find a way to mitigate that incentive and incur the costs of implementing the chosen solution. Thus, the costs of administering and enforcing this rule are not a consequence of the proposed non-discrimination rule as such, but of the decision to allow charging for Quality of Service, and are therefore better discussed in the context of that decision.[538]

---

[534] On the importance of certainty for network providers and application developers, see footnotes 226 to 231 above and accompanying text.

[535] For example, neither the FCC's Internet Policy Statement nor the FCC's Notice of Proposed Rulemaking in the Open Internet Proceeding further defined the term "reasonable." Federal Communications Commission (2005); Federal Communications Commission (2009b), para 135.

[536] These exceptions are afflicted with the same problems as non-discrimination rules that use ambiguous or undefined terms to describe which discriminatory behavior is banned. See the analysis in Sections "Ban Discrimination That is Unreasonable" and "Problems with Case-by-Case Adjudications" above.

[537] See footnote 479 above and accompanying text.

[538] A network neutrality regime could avoid this problem by allowing network providers to offer Quality of Service, but prohibiting them from charging for it. This solution has its own costs and benefits. See van Schewick (2010a), pp.10-11. Like all questions related to charging for Quality of Service, this question is outside the scope of this paper.

In sum, the rule restricts the evolution of the network to some degree, but only to the extent necessary to realize the goals of network neutrality regulation. The costs of administering and enforcing the non-discrimination rule are considerably lower than those of most of the other proposed non-discrimination rules. And while the rule reduces network providers' profits and, potentially, incentives to invest in more and better broadband networks to some degree by preventing network providers from freely engaging in discriminatory conduct and from charging application and content providers for Quality of Service-enhanced access to their Internet service customers, it allows network providers to profit in ways (for example, by charging end users for Quality of Service, or by engaging in application-agnostic forms of price discrimination) that other proposals would forbid. Thus, the rule does not impose more social costs than necessary to protect the values that network neutrality regulation is designed to protect.

THE OPEN INTERNET ORDER'S NON-DISCRIMINATION RULE

In October 2009, the Federal Communications Commission started the Open Internet Proceeding by publishing a Notice of Proposed Rulemaking that proposed a set of network neutrality rules and asked for comment. The proposed rules included a strict non-discrimination rule that would have required network providers to treat every packet the same, subject to reasonable network management.[539]

Over the course of the proceeding, the FCC held several public workshops, organized a technical advisory process and received more than 100,000 written comments.[540] All of the proposals for non-discrimination rules discussed in this paper were supported by some commenters.[541]

In December 2010, FCC Chairman Julius Genachowski circulated draft Open Internet Rules to the other commissioners. The proposed rules were based on a proposal for a network neutrality bill that had been negotiated by Rep. Henry A. Waxman, the Chairman of the House Committee on Energy and Commerce, and Rep. Rick Boucher, the Chairman of the House Subcommittee on Communications, Technology and the Internet, with the large phone and cable network providers, Internet companies, consumer groups and open Internet groups in the Fall of 2010.[542] Not all participants in the negotiations supported the final proposal.[543] The proposal was never introduced. Rep. Waxman and Rep. Boucher had abandoned the effort when the Republican members of the Energy and Commerce Committee refused to support the bill.[544]

The rules proposed by the Chairman included a non-discrimination rule that applied to fixed, but not to mobile Internet access service. Like the non-discrimination rule in the Waxman proposal, the rule banned "unreasonable" discrimination without specifying how to interpret the

[539] See footnote 101 above.

[540] Federal Communications Commission (2010c), p.2, para 2.

[541] See the sources cited in support of the different proposals above.

[542] Hart (2010b). In a number of meetings and conversations with FCC officials in November 2010, AT&T representatives argued the Waxman bill should be a model for the FCC's Open Internet Rules (Shields (2010a); AT&T (2010b)). The Open Internet Rules regarding blocking, non-discrimination and disclosure as well as the definitions of broadband Internet access service and of reasonable network management adopted by the FCC in December 2010 closely follow the corresponding provisions of the Waxman bill (compare Federal Communications Commission (2010c), pp. 88-89 and Draft Bill To Amend Title I of the Communications Act of 1934 to Provide for Internet Openness, and for Other Purposes (2010)). For more on the Waxman proposal, see footnote 223 above.

[543] In particular, Free Press and the Open Internet Coalition did not support the compromise proposal. Jerome (2010a). Eggerton (2010). Even consumer groups, organizations and firms that had supported the Waxman proposal as a legislative measure during the negotiations in Congress in September 2010 later opposed adopting it as agency rules. See the discussion in Public Knowledge & Media Access Project (2010) (disclosing a meeting in which representatives of Consumers Union, Free Press, Media Access Project, New America Foundation, Public Knowledge, Writers Guild West, Amazon.com, Dish Network, Netflix and Skype expressed their opposition to adopting the Waxman proposal as agency rules).

[544] Hart (2010a); Waxman (2010).

term and left it to later case-by-case adjudication to decide whether specific discriminatory conduct meets this criterion.[545]

While the strategic interests of regulators or legislators considering the adoption of network neutrality rules and of the big stakeholders on both sides of the network neutrality debate are aligned in favor of such a rule,[546] that rule does not adequately protect the values that network neutrality regulations are designed to protect.[547] Thus, from the perspective of network neutrality proponents, this proposal was highly unsatisfactory.[548]

The two Republican commissioners, Commissioner Robert M. McDowell and Commissioner Meredith Atwell Baker, had made clear that they would not support any network neutrality rules, so the Chairman needed the votes of the two other Democratic commissioners, Commissioner Michael J. Copps and Commissioner Mignon L. Clyburn. Both supported considerably stronger network neutrality rules than the ones that the Chairman was proposing.[549] Commissioner Copps publicly threatened to vote against the draft rules in the form proposed by the Chairman. A "no" from either of them would have killed the proposal.[550]

This put the Chairman in a difficult position: During his presidential campaign, President Obama had promised to enact network neutrality rules.[551] Attempts to enact network neutrality rules in Congress had failed.[552] The Democrats had lost the House in November 2010, and with the Republicans in the majority, it was clear that the House would not support any network-neutrality-related action in the future.[553]

At the same time, the Chairman felt he needed AT&T's support to deflect Republican criticism in Congress in the future, and AT&T strongly supported using the Waxman proposal as a basis for the Commission's network neutrality rules.[554,555]

---

[545] The draft rules were not released publicly, but described by the Chairman in public remarks when he circulated the draft rules: "And so the proposed framework includes a bar on unreasonable discrimination in transmitting lawful network traffic." Genachowski (2010). For the non-discrimination rule in Waxman's proposal, see Draft Bill To Amend Title I of the Communications Act of 1934 to Provide for Internet Openness, and for Other Purposes (2010), Section 12(a)(1)(B). The non-discrimination rule in the Waxman proposal applied to wireline broadband Internet access service and would have left the decision whether to treat fixed wireless and satellite Internet access service like wireline or wireless Internet access to the FCC. See Draft Bill To Amend Title I of the Communications Act of 1934 to Provide for Internet Openness, and for Other Purposes (2010), Section 12(c).

[546] See footnotes 273 to 281 above and accompanying text.

[547] See Sections "Ban Discrimination That Is Unreasonable" and "Problems with Case-by-Case Adjudication" above.

[548] Network neutrality proponents also opposed other aspects of the proposal. For an overview, see van Schewick (2010l); van Schewick (2010b).

[549] For a summary of their criticisms, see Hatch (2010); Copps (2010); Clyburn (2010).

[550] Hatch (2010).

[551] Obama for America (2007).

[552] The last such attempt had failed in September 2010, when the Republican leadership in Congress refused to support the compromise proposal negotiated by Rep. Waxman and Rep. Boucher. Hart (2010a); Waxman (2010).

[553] Hart (2010b); The Hill Administrator (2010).

[554] At the time, AT&T was the second-largest provider of wireline Internet access in the US, the largest DSL provider in the US and the second largest wireless provider in the US. Leichtman Research Group (2010). AT&T, which has consistently outspent every other corporation in donations to congressional campaigns in the past, has considerable

The Chairman and the two Democratic commissioners negotiated over improvements to the order up to the day before the Commission's Open Meeting on December 21.[556] In the end, they agreed on the following compromise:[557] They did not change the text of the non-discrimination rule. They agreed, however, to change the text of the order to provide more clarity to industry participants and to provide guidance to future adjudications.

The Open Internet Rules were adopted at the FCC's Open Meeting in December 2010 and went into effect in November 2011. The rules were appealed by Verizon, Free Press and others and vacated by the Court of Appeals for the D.C. Circuit in January 2014.[558]

In the final Open Internet Report and Order, the FCC adopted a non-discrimination rule that bans providers of fixed broadband Internet access service from "unreasonably discriminat[ing] in transmitting lawful network traffic over a consumer's broadband Internet access service. Reasonable network management shall not constitute unreasonable discrimination."[559] Whether certain discriminatory conduct qualifies as unreasonable will be determined by the FCC in case-by-case adjudications.

The text of the order specifies the factors that the FCC will use to determine whether a certain discriminatory conduct constitutes unreasonable discrimination:[560] transparency (i.e. whether differential treatment is disclosed), end-user control and end-user choice, use-agnostic discrimination and conformity of the practice with "best practices and technical standards adopted by open, broadly representative, and independent Internet engineering, governance initiatives, or standards-setting organizations."[561]

---

influence in Washington. See, e.g., Mayer (2009); Lasar (2009); Lasar (2010); New York Times (2010); Kang (2009) (discussing data showing that 70 of the 72 democratic members of congress who sent a letter to Chairman Genachowski in October 2009 cautioning against adopting network neutrality rules had received contributions from Internet service providers, with AT&T donating the most, and donating to 58 out of the 72).

[555] The Hill Administrator (2010) ("Public advocates are concerned about how much Genachowski appears to be listening to AT&T, with one saying he has practically given them 'veto powers.'"); Shields (2010a) (discussing a large number of conversations between Jim Cicconi, Senior Executive Vice President External & Legislative Affairs at AT&T and head of AT&T's public policy organization, and Edward Lazarus, the Chairman's Chief of Staff, in November 2011 and citing a number of network neutrality proponents interpreting the Chairman's efforts to get AT&T's approval as an attempt to minimize opposition against the rules and protect the Commission against Republican criticism); AT&T (2010b) (disclosing two conversations between Cicconi and Lazarus, in which "Mr. Cicconi discussed the merits of the proposed Waxman legislation and why it should be a model for a substantive resolution of the issues raised in these proceedings.").

[556] Gross (2010); Hesseldahl (2010).

[557] This paper focuses on non-discrimination rules. For an overview of the other parts of the compromise, see van Schewick (2010c).

[558] Federal Communications Commission (2011a), pp. 59192, 59223; Federal Communications Commission (2011b)); Verizon v. FCC, No. 11-1355 (D.C.Cir. 2014), available online at http://www.cadc.uscourts.gov/internet/opinions.nsf/3AF8B4D938CDEEA685257C6000532062/$file/11-1355-1474943.pdf.

[559] Federal Communications Commission (2010c), p. 88, Section 8.7.

[560] Federal Communications Commission (2010c), pp. 42-42, paras 69-74. The following paragraphs are based in part on van Schewick (2010c).

[561] Federal Communications Commission (2010c), p. 42, para 74.

Use-agnostic discrimination (or "application-agnostic" discrimination), the order explains, is "differential treatment that does not discriminate among specific uses of the network or classes of uses." According to the order, use-agnostic discrimination is likely to be reasonable, which suggests, in turn, that differential treatment that discriminates among specific uses of the network or classes of uses is likely to be unreasonable.[562] This is the same substantive standard as the one used by the non-discrimination rule proposed by this paper (see *Table 2: The Open Internet Order's Non-Discrimination Rule and the Rule Proposed by this Paper* below). As explained above, allowing use-agnostic discrimination, but banning discrimination among uses or classes of uses preserves the application-agnosticism of the network.[563]

Table 2
The Open Internet Order's Non-Discrimination Rule and the Rule Proposed by this Paper

	Application-Agnostic Discrimination	**Application-Specific Discrimination**
FCC's Non-Discrimination Rule	likely to be reasonable	likely to be unreasonable
Non-Discrimination Rule Proposed by this Paper	allowed	banned

Thus, in evaluating whether discriminatory conduct is reasonable, the FCC will consider how well the conduct preserves two of the four factors – application-agnosticism and user choice – that have fostered application-innovation and allowed the Internet to serve as a platform for social, political and cultural interaction in the past.[564] As Commissioner Copps explained in his concurring statement, this was a deliberate decision: "In discussing the 'no unreasonable discrimination' standard, we put particular emphasis on keeping control in the hands of users and preserving an application-blind network—a key part of making the Internet the innovative platform it is today."[565] In addition, the first section of the Open Internet Rules lists preserving end-user control, end-user choice and the freedom to innovate without permission

[562] Federal Communications Commission (2010c), p. 41-42, para 73.

[563] See footnote 454 and accompanying text and *Box 22: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination* above.

[564] On these factors (the other two are innovation without permission and low cost of application-innovation), see footnotes 53 to 56 and accompanying text, *Box 4: Application-Agnostic v. Application-Blind* and *Box 3: The Importance of User Choice* above.

[565] Copps (2010), p. 3. See also Federal Communications Commission (2010c), pp. 40-41, para 71: "Maximizing end-user control is a policy goal Congress recognized in Section 230(b) of the Communications Act, and end-user choice and control are touchstones in evaluating the reasonableness of discrimination. As one commenter observes, 'letting users choose how they want to use the network enables them to use the Internet in a way that creates more value for them (and for society) than if network providers made this choice,' and 'is an important part of the mechanism that produces innovation under uncertainty.'" (ibid., references omitted); Federal Communications Commission (2010c), p. 42, para 73: "Use-agnostic discrimination (sometimes referred to as application-agnostic discrimination) is consistent with Internet openness because it does not interfere with end users' choices about which content, applications, services, or devices to use. Nor does it distort competition among edge providers." (ibid., references omitted).

(the third of the four factors that have fostered application-innovation in the past) as explicit purposes of the rules.[566]

As this paper has shown, using these factors – application-agnosticism, user choice and innovation without permission[567] – as guidelines for evaluating behavior provides clear answers regarding which types of discriminatory behavior should and should not be allowed. For example, while the order does not discuss how the different forms of Quality of Service discussed in this paper would be evaluated under the "no unreasonable discrimination" standard, the analysis in this paper suggests which results the FCC will reach, if it takes these factors seriously (see *Table 3: Evaluating Different Forms of Quality of Service Under the Open Internet Order's Non-Discrimination Rule* below). In particular, under the FCC's standard as clarified by the order, user-controlled Quality of Service where (1) the different classes of service are offered equally to all applications and classes of applications, (2) the user is able to choose whether and when to use which class of service, and (3) the network provider is allowed to charge only its own Internet service customers for the use of the different classes of service,[568] is likely to be reasonable.[569,570] By contrast, offering Quality of Service exclusively to one or more applications within a class of "like" applications is unlikely to be reasonable.[571] Offering different types of service to different provider-defined classes of applications is also likely to be unreasonable, even if the network provider treats like traffic alike (that is, even if it does not

[566] 47 C.F.R. §8.1: "The purpose of this Part is to preserve the Internet as an open platform enabling consumer choice, freedom of expression, end-user control, competition, and the freedom to innovate without permission." For a discussion of the importance of innovation without permission in the order, see Federal Communications Commission (2010c), p. 3, para 3,pp. 5-6, para 13.

[567] Since preserving the freedom to innovate without permission is an explicit purpose of the Open Internet Rules, this factor can be used to interpret any provision of the rules, including the non-discrimination rule.

[568] According to the Open Internet Order, charging application or content providers who are not the network provider's Internet access service customers for prioritized or otherwise enhanced access to its Internet access service customers is unlikely to be reasonable under the order's non-discrimination rule. While the Open Internet Rules do not directly address this question, the text of the order discusses it in detail. The order explicitly endorses the concerns against these arrangements ((Federal Communications Commission (2010c)), paras 76 and 24-34), unequivocally rejects the main arguments in favor of them (ibid, paras 40 and 28), and concludes that "as a general matter," arrangements of this kind are "unlikely" to be considered reasonable under the non-discrimination rule (ibid, para 76). Thus, under the Open Internet Rules, network providers are (most likely) allowed to charge only their own Internet access service customers for any differential treatment allowed by the non-discrimination rule. For a more detailed analysis, see van Schewick (2010c), Section 2.

[569] See *Table 3: Evaluating Different Forms of Quality of Service Under the Open Internet Order's Non-Discrimination Rule* below and footnotes 471 to 483 and accompanying text above.

[570] User-controlled Quality of Service is the only type of Quality of Service that the order discusses explicitly: "Thus, enabling end users to choose among different broadband offerings based on such factors as assured data rates and reliability, or to select quality-of-service enhancements on their own connections for traffic of their choosing, would be unlikely to violate the no unreasonable discrimination rule, provided the broadband provider's offerings were fully disclosed and were not harmful to competition or end users." Federal Communications Commission (2010c), p. 41, para 71.

[571] See *Table 3: Evaluating Different Forms of Quality of Service Under the Open Internet Order's Non-Discrimination Rule* below and footnotes 368 to 378 and accompanying text above.

discriminate among classes of applications that are alike and does not discriminate among applications within a class of like applications).[572]

In addition to specifying which factors should be used in evaluating discriminatory conduct under the FCC's "no unreasonable discrimination" standard, the order explicitly rejects some alternative interpretations. Some commenters had supported using an antitrust framework to distinguish socially beneficial from socially harmful discrimination.[573] The order explicitly rejects the view that the non-discrimination rule should only prohibit discrimination that is "anticompetitive."[574] As explained above, such a rule (or an interpretation of the FCC's rule that restricted unreasonable discrimination to discrimination that is anticompetitive) would have made it impossible to bring complaints against many types of discriminatory conduct that network neutrality proponents are concerned about.[575]

Finally, according to the order, the same principles that guide the Commission's interpretation of the non-discrimination rule will also guide the Commission's evaluation of network management practices under the Open Internet Rules' exception for reasonable network management.[576] The exception applies to the no-blocking rules for fixed and mobile Internet access and to the non-discrimination rule. Some had argued that discriminatory and exclusionary practices should automatically qualify as "reasonable network management," as long as they were designed to solve network management problems.[577] However, the harm to users and innovators from discriminatory or exclusionary conduct is the same regardless of the network provider's motivation, making it necessary to impose stronger constraints on reasonable network management.[578] In line with these considerations, the order makes clear that network management will be evaluated by the same principles that guide the interpretation of the non-discrimination rule.

Overall, the non-discrimination rule adopted by the Commission (as clarified by the text of the order) constitutes a considerable improvement over the same rule without clarifications.

---

[572] See *Table 3: Evaluating Different Forms of Quality of Service Under the Open Internet Order's Non-Discrimination Rule* below and footnotes 380 to 436 and accompanying text above.

[573] See the sources cited throughout Section "Ban Discrimination That Violates an Antitrust Framework" above.

[574] Federal Communications Commission (2010c), pp. 45-46, para 78: "We also reject the argument that only "anticompetitive" discrimination yielding "substantial consumer harm" should be prohibited by our rules. We are persuaded those proposed limiting terms are unduly narrow and could allow discriminatory conduct that is contrary to the public interest. The broad purposes of this rule—to encourage competition and remove impediments to infrastructure investment while protecting consumer choice, free expression, end-user control, and the ability to innovate without permission—cannot be achieved by preventing only those practices that are demonstrably anticompetitive or harmful to consumers. Rather, the rule rests on the general proposition that broadband providers should not pick winners and losers on the Internet—even for reasons that may be independent of providers' competitive interests or that may not immediately or demonstrably cause substantial consumer harm." (references omitted)

[575] See the discussion in Section "Ban Discrimination That Violates an Antitrust Framework" above.

[576] Federal Communications Commission (2010c), p. 50, para 87.

[577] AT&T (2010a), pp. 183-187; Brito, et al. (2010), pp. 22-23; Cox Communications (2010), pp. 30-33.

[578] See, e.g., van Schewick (2008a), pp. 4-8. See also footnotes (and accompanying text) 207 to 210, 245 to 252 and 262 to 269 above.

The general theoretical framework underlying the order as well as the specific factors that will be used to interpret the non-discrimination rule and the reasonable-network-management exception are in line with the broader theoretical framework that calls for network neutrality regulation are based on.[579] Contrary to the standards used by other case-by-case approaches (e.g., by an antitrust framework or the Google-Verizon legislative framework), the factors the FCC will use to evaluate differential treatment do not automatically exclude instances of discrimination that network neutrality proponents are concerned about.[580] Instead, the substantive factors – application-agnosticism and user choice – reinforce key values that were at the core of the Internet's success. Specifying the factors provides additional clarity to market participants and guidance to the bureaus within the FCC which may end up enforcing the order.

Still, compared with the bright-line non-discrimination rule supported by this paper, considerable uncertainty remains. For example, it is not clear how the Commission would evaluate a practice that is in line with some, but not all factors mentioned by the order. Thus, the rule leaves a lot of discretion to later adjudicators. By creating considerable uncertainty that will only be resolved in later case-by-case adjudications, the rule creates many of the same problems and social costs as the less precise standards discussed above.[581]

The text of the order provides the most certainty with respect to behavior that is likely to be reasonable. In particular, the order explains that differential treatment that is use-agnostic is likely to be reasonable and that "end-user choice and control are touchstones of evaluating reasonableness."[582] This suggests that network providers who would like to minimize the risk of having to defend themselves in costly and highly public adjudications at the FCC should choose practices and invest in network technologies that are use-agnostic (i.e. that do not discriminate among specific uses or classes of uses) and preserve user choice over technologies and practices that do not.

[579] See the discussion in Section "Criteria for Evaluating Non-Discrimination Rules" above.

[580] See Sections "Ban Discrimination that Violates an Antitrust Framework" and "Ban Discrimination that is Anticompetitive or Harms Users" above.

[581] For a discussion of these costs, see Section "Problems with Case-by-Case Adjudication."

[582] Federal Communications Commission (2010c), p. 40-42, paras 71-73.

Table 3
Evaluating Different Forms of Quality of Service Under the Open Internet Order's Non-Discrimination Rule

Types of QoS / Criteria for evaluation	Provider-controlled QoS to individual applications within a class of like applications	Provider-controlled QoS to provider-defined classes of applications	User-controlled QoS where (1) the different classes of service are offered equally to all applications and classes of applications; (2) the user is able to choose whether and when to use which class of service; (3) the network provider charges only its own Internet service customers for the use of the different classes of service.*
Preserves application-agnosticism of the network/is use-agnostic (i.e. does not discriminate among specific uses or classes of uses)	No	No	Yes
Preserves user choice	No	No	Yes
Preserves innovation without permission	No	No	Yes
Likely to be reasonable under the FCC's non-discrimination rule	No	No	Yes

* According the FCC's Open Internet Order, charging application or content providers who are not the network provider's Internet access service customers for prioritized or otherwise enhanced access to the network provider's Internet access service customers is likely to be unreasonable. Thus, under the Open Internet Rules, network providers are allowed to charge only their own Internet access service customers for any differential treatment allowed by the non-discrimination rule. See footnote 568.

CONCLUSION: NETWORK NEUTRALITY AND QUALITY OF SERVICE

The network neutrality debate is often framed as a debate for or against Quality of Service. As the paper shows, the reality is much more nuanced. Some proposals take an all-or-nothing approach to discrimination. They ban or allow all forms of discrimination and, consequently, Quality of Service. Most proposals take a more nuanced position. They allow some, but not all forms of Quality of Service, with different proposals drawing the line between acceptable and unacceptable forms of Quality of Service in different ways.

Often, it is not immediately apparent how a specific non-discrimination rule affects network providers' ability to offer Quality of Service. To address this problem, this paper explores the effect of the various proposals on the different forms of Quality of Service. The results of this analysis are summarized in *Table 4: The Impact of Different Non-Discrimination Rules on Quality of Service* below.

Underlying the differences between the proposals are disagreements over the social benefits and costs of the different forms of Quality of Service. In this respect, the paper offers interesting new insights.

Most network neutrality proponents agree that *allowing network providers to offer Quality of Service exclusively to one or more applications within a class of "like" applications* should be prohibited and this paper shares that view.[583] For example, a network provider should not be allowed to offer a low-delay service only to its own Internet video application, or only to select unaffiliated Internet video applications. This type of Quality of Service interferes with users' ability to use the applications of their choice without interference from network providers and enables network providers to use the provision of Quality of Service as a tool to distort competition among applications within a class, which is exactly what network neutrality rules are designed to prevent.

By contrast, many network neutrality proponents see no problems *with allowing network providers to offer different types of service to different provider-defined classes of applications, as long as the network provider treats like traffic alike*. In other words, they would allow network providers to provide different types of service to different provider-defined classes of applications that are not alike, as long as they do not discriminate among classes of applications that are alike or among applications within a class of like applications. This requirement is often called "like treatment."[584] Under this approach, a network provider would be allowed to offer low-delay service to Internet telephony, but not to e-mail, as long as it does not treat Vonage

[583] On this form of Quality of Service, see Section "Ban Discrimination Among Like Applications and Classes of Applications" in the paper.

[584] On this form of Quality of Service, see Section "Allow Discrimination Among Classes of Applications That Are Not Alike" in the paper.

differently from Skype, or Gmail differently from Hotmail.[585] In the US, the AT&T BellSouth Merger conditions and various draft bills in Congress allowed this form of Quality of Service.

The positive stance towards forms of Quality of Service that provide like treatment is based on the assumption that discriminating among classes of applications that are not alike is socially harmless and should therefore be allowed. As this paper shows, this assumption is not correct. In many cases, discrimination among classes of applications hurts some classes of applications, even if the classes are not alike. For example, some Internet applications such as Internet telephony applications, Internet messaging applications or Internet video offerings compete with network-provider services that are sold separately from Internet access and do not run over the Internet-access portion of the network provider's access network. In these cases, discriminating against all applications in that class allows the network provider to favor its own offering without discriminating among applications within the class. Moreover, applications in a class can be harmed by differential treatment even if they do not compete directly with applications in other classes that are treated more favorably.

In addition, like treatment negatively affects several of the factors that have fostered application innovation in the past. *First*, like treatment removes the *application-agnosticism* of the network. Allowing network providers to treat classes of applications differently requires the network provider to identify the different applications on its network in order to decide which class they belong to and determine the appropriate type of service. Thus, like treatment requires network providers to treat data packets differently based on information about the applications on the network. Since the concept of "like applications" is not well defined, network providers have broad discretion to decide which applications are alike, which allows them to deliberately or inadvertently distort competition among applications or classes of applications. *Second*, like treatment violates the principle of *user choice*. Under like treatment, network providers, not users, choose which application should get which Quality of Service. Since users' preferences for Quality of Service are not necessarily the same across users and may even vary for the same user over time, letting network providers determine which applications gets which Quality of Service will result in levels of Quality of Service that do not meet users' needs. *Third*, like treatment harms application innovation by requiring innovators to convince network providers that their application belongs to a certain class. Requiring network providers to take action before an application can get the Quality of Service it needs violates the principle of *innovation without permission* and reduces the chance that new applications actually get the type of service they need. *Finally,* disputes over which classes of applications are alike, or whether a certain application belongs to a certain class, are likely to be frequent and difficult to resolve, creating high costs of regulation.

Thus, contrary to what is commonly assumed, forms of Quality of Service that respect the principle of like treatment do not adequately protect the values that network neutrality is designed to protect and should not be allowed under a network neutrality regime.

[585] Internet telephony is sensitive to delay, but e-mail is not, so the two classes of applications are not alike. See footnote 10 above.

By contrast, Quality of Service architectures *where network providers make different types of service available equally to all applications and classes of applications and where users choose whether and when to use which type of service* do not raise similar concerns.[586] *First*, they preserve the *application-agnosticism* of the network: The provision of Quality of Service is not dependent on which applications users are using, but on the Quality-of-Service-related choices that users make; thus, the network providers does not need to know anything about which applications are using its network in order for this scheme to work. The network provider only makes different classes of service available, but does not have any role in deciding which application gets which Quality of Service; this choice is for users to make. As a result, network providers cannot use the provision of Quality of Service as a mechanism to distort competition among applications or classes of applications. *Second*, since users choose when and for which applications to use which type of service (in line with the principle of *user choice*), they can get exactly the Quality of Service that meets their preferences, even if these preferences differ across users or (for a single user) over time. *Third*, in line with the principle of *innovation without permission*, an innovator does not need support from the network provider in order for his application to get the Quality of Service it needs. The only actors who need to be convinced that the application needs Quality of Service are the innovator, who needs to communicate this to the user, and the user, who wants to use the application. This greatly increases the chance that an application can get the type of service it needs.

In sum, this type of user-controlled Quality of Service offers the same potential social benefits as other, discriminatory or provider-controlled forms of Quality of Service without the social costs. With appropriate restrictions on charging and with provisions that protect the quality of the baseline service from dropping below unacceptable levels, this type of Quality of Service should be allowed under a network neutrality regime. Under the non-discrimination rule proposed by this paper and the FCC's non-discrimination rule, these are the only forms of quality of Service that network providers would be able to offer.

If policy makers adopt a non-discrimination rule that allows network providers to offer some form of Quality of Service, they need to decide whether, and if so, whom network providers should be allowed to charge for it. Again, policy makers have a number of options, each supported by at least some proponents of network neutrality: (1) the network provider is not allowed to charge anyone for the use of Quality of Service (though it can increase the general price for Internet service); (2) it can charge only its Internet service customers; (3) it can charge its Internet service customers and/or application and content providers, but is required to offer the service to application and content providers on a non-discriminatory basis; (4) it can charge its Internet service customers and/or application and content providers. Concerns about offering differential treatment and about charging for it are driven by different sets of policy considerations which should be considered and evaluated separately. I take up restrictions on

[586] On this type of Quality of Service, see Section "Ban Application-Specific Discrimination, Allow Application-Agnostic Discrimination," Subsection "Allowing the Network to Evolve" in the paper.

charging elsewhere.[587] There, I argue that network providers should only be allowed to charge their own Internet service customers for any differential treatment allowed by the non-discrimination rule.[588]

Finally, a network provider who is allowed to charge for Quality of Service has an incentive to degrade the quality of the baseline, best-effort service to motivate users to pay for an enhanced type of service. To mitigate this problem, any network neutrality regime that allows network providers to charge for Quality of Service should require the regulatory agency in charge of enforcing the network neutrality rules to monitor the quality of the baseline service and set minimum quality standards, if the quality of the baseline service drops below appropriate levels.[589]

Opponents of network neutrality regulation have created the impression that policy makers need to choose between protecting users and application innovators against interference from network providers on the one hand and innovation in the network and the needs of network providers on the other hand. As the paper shows, it is possible to protect users and innovators while giving network providers the tools they need to manage their networks and allowing the network to evolve. Thus, regulators can have their cake and eat it, too.

[587] van Schewick (2010h); van Schewick (2010i); van Schewick (2010d), pp. 278-280, 290-293

[588] This restriction would not constrain interconnection agreements in any way. Thus, payments among interconnecting networks would remain possible.

[589] For a discussion of this requirement, see van Schewick (2010a), pp. 10-11. The European Union has adopted a similar rule following its review of the regulatory framework for telecommunications services. See Article 22(3) of the Universal Service Directive; European Commission (2007), pp. 92, 95-97, 101.

Table 4

The Impact of Different Non-Discrimination Rules on Quality of Service

Forms of QoS / Rules	Provider-controlled QoS to individual applications within a class of like applications	QoS to provider-defined classes of applications	User-controlled QoS
Allow all discrimination	allowed*	allowed*	allowed*
Case-by-case approaches	unclear	unclear	unclear
Ban discrimination that is not disclosed	allowed, if disclosed*	allowed, if disclosed*	allowed, if disclosed*
Ban discrimination among like applications and classes of applications, allow discrimination among classes of applications that are not alike and application-agnostic discrimination ("like treatment")	banned	allowed, as long as like traffic is treated alike*	allowed, if (1) different classes of service are made available equally to all applications and classes of applications; AND (2) user is able to choose whether and when to use which class of service.*
Ban application-specific discrimination, allow application-agnostic discrimination	Banned	banned	allowed, if (1) different classes of service are made available equally to all applications and classes of applications; (2) user is able to choose whether and when to use which class of service.*
FCC's non-discrimination rule	likely banned	likely banned	likely allowed, if (1) different classes of service are made available equally to all applications and classes of applications; (2) user is able to choose whether and when to use which class of service; AND (3) network provider charges only its own Internet service customers for the use of the different classes of service.**
Ban all discrimination	Banned	banned	banned

* If policy makers adopt a non-discrimination rule that allows network providers to offer some form of Quality of Service, they need to decide whether, and if so, whom network providers should be allowed to charge for it. As I argue elsewhere, network providers should only be allowed to charge their own Internet service customers for any differential treatment allowed by the non-discrimination rule.

** According the FCC's Open Internet Order, charging application or content providers who are not the network provider's Internet access service customers for prioritized or otherwise enhanced access to its Internet access service customers is likely to be unreasonable. Thus, under the Open Internet Rules, network providers are allowed to charge only their own Internet access service customers for any differential treatment allowed by the non-discrimination rule.

REFERENCES

Akhund, Immad, Brian Ascher, Aneel Bhusri, Matt Blumberg, Brad Burnham, Stewart Butterfield, et al. 2010. *Letter to Chairman Genachowski.* Appendix B to the Open Internet Coalition's Comments In the Matter of Preserving the Open Internet submitted January 14, 2010. GN Dkt. No. 09-191. http://apps.fcc.gov/ecfs/document/view?id=7020377928

Al Jazeera & Reuters. 2014. "Net Neutrality Wins in Brazil's 'Internet Constitution'." *Al Jazeera America.* March 26. http://america.aljazeera.com/articles/2014/3/26/brazil-internet-constitution.html

Allot Communications & Openet. 2010. *Managing the Unmanageable: Monetizing and Controlling OTT applications. FierceLive! Webinar Presentation.* Attachment to Free Pree's Ex Parte Letter In the Matter of Preserving the Open Internet submitted December 14, 2010. GN Dkt. No. 09-191. http://apps.fcc.gov/ecfs/document/view?id=7020923750

Ammori, Marvin. 2009. "Net Neutrality and the 21st Century First Amendment." *Balkinization.* December 10. http://balkin.blogspot.com/2009/12/net-neutrality-and-21st-century-first.html

Ammori, Marvin. 2010a. "All the Political Cover the FCC Could Ever Want." *The Huffington Post.* September 30. http://www.huffingtonpost.com/marvin-ammori/all-the-political-cover-t_b_745243.html

Ammori, Marvin. 2010b. "A Guide to the Network Neutrality Discussions at the FCC." *The Huffington Post.* August 4. http://www.huffingtonpost.com/marvin-ammori/a-guide-to-the-network-ne_b_670784.html

Anderson, Nate. 2007. "Deep Packet Inspection Meets 'Net Neutrality, CALEA." *Ars Technica.* July 25. http://arstechnica.com/articles/culture/Deep-packet-inspection-meets-net-neutrality.ars

Areeda, Phillip E. & Herbert Hovenkamp. 2010. *Antitrust Law: An Analysis of Antitrust Principles and Their Application* 3rd ed. Vol. VIII. New York, NY: Aspen Publishers.

Areeda, Phillip E. & Herbert Hovenkamp. 2008. *Antitrust Law: An Analysis of Antitrust Principles and Their Application.* 3rd ed. Vol. III New York, NY: Aspen Publishers.

Areeda, Phillip E., Herbert Hovenkamp & John L. Solow. 2007. *Antitrust Law: An Analysis of Antitrust Principles and Their Application* Vol. IIB. New York, NY: Aspen Publishers.

Armitage, Grenville J. 2003. "Revisiting IP QoS: Why Do We Care, What Have We Learned? ACM SIGCOMM 2003 RIPQOS Workshop Report." *Computer Communication Review*, 33(5): 81-88.

AT&T. 2006. *Merger Commitments.* Attachment to AT&T's Ex Parte Communication In the Matter of AT&T and BellSouth Applications for Approval of Transfer submitted December 28, 2006. WC Dkt. No. 06-74. http://apps.fcc.gov/ecfs/document/view?id=6518716381

AT&T. 2009. AT&T Response to Wireless Telecommunications Bureau Letter. DA 09-1737 (July 31, 2009), Letter to Federal Communications Commission. WT Dkt. No. RM-11361. August 21. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020036306

AT&T. 2010b. Notice of Ex Parte Conversations to Federal Communications Commission. GN Dkt. No. 09-191. November 22. http://apps.fcc.gov/ecfs/document/view?id=7020921306

AT&T, Alcatel-Lucent, Cable Europe, Ciena, Cisco, Ericsson, et al. 2008. *Ensuring Network Stability and Consumer Confidence in Competitive Markets.* http://www.laquadrature.net/files/net.confidence.coalition.pdf

AT&T. 2007. Comments to Federal Communications Commission. WC Dkt. No. 07-52. June 15. http://apps.fcc.gov/ecfs/document/view?id=6519529324

Atkinson, Robert D. & Philip J. Weiser. 2006. "A Third Way on Network Neutrality." *The New Atlantis*(Number 13, Summer 2006): 47-60.

Bachula, Gary R. 2006. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: Net Neutrality.* http://commerce.senate.gov/pdf/bachula-020706.pdf

Bagley, Nicholas & Richard L. Revesz. 2006. "Centralized Oversight of the Regulatory State." *Columbia Law Review*, 106(6): 1260-1329.

Balkin, Jack, John Blevins, Jim Chen, Lawrence Lessig, Barbara van Schewick & Tim Wu. 2009. Ex Parte Letter to Federal Communications Commission. GN Dkt. No. 09-51. November 2. http://apps.fcc.gov/ecfs/document/view?id=7020244236

Balkin, Jack M. 2008. "The Future of Free Expression in a Digital Age." *Pepperdine Law Review*, 36(1).

Balkin, Jack M. 2009. *Testimony before the Federal Communications Commision at its Workshop on Speech, Democratic Engagement, and the Open Internet.* Federal Communications Commission. http://www.fcc.gov/events/workshop-speech-democratic-engagement-and-open-internet

Balkin, Jack M. , Jim Chen, Lawrence Lessig, Barbara van Schewick & Tim Wu. 2009. Brief Amicus Curiae Urging That the FCC's Order Be Affirmed to United States Court of Appeals for the District of Columbia Circuit. No. 08-1291. October 5. http://www.law.stanford.edu/display/images/dynamic/publications_pdf/van%20Schewick%20Amicus%20Brief.pdf

Bar-Gill, Oren. 2006. "Bundling and Consumer Misperception." *University of Chicago Law Review*, 73(Winter): 33-61.

Barkow, Rachel E. 2010. "Insulating Agencies: Avoiding Capture Through Institutional Design." *Texas Law Review*, 89(1): 15-79.

Bastian, C., T. Klieber, J. Livingood, J. Mills & R. Woundy. 2010. "Comcast's Protocol-Agnostic Congestion Management System." Request for Comments 6057. IETF.

Bauer, Steven, David Clark & William Lehr. 2009. "The Evolution of Internet Congestion." Paper presented at 37th Research Conference on Communication, Information and Internet Policy (TPRC 2009). Arlington, VA, USA.

Baumol, William J., Martin Cave, Peter Cramton, Robert Hahn, Thomas W. Hazlett, Paul L. Joskow, et al. 2007. *Economists' Statement on Network Neutrality Policy.* AEI-Brookings Joint Center Working Paper RP07-08. http://ssrn.com/abstract=976889

BBC News. 2012. "Social Messaging Apps 'Lost Networks $13.9bn in 2011'." *BBC News* February 21. http://www.bbc.co.uk/news/technology-17111044

Becker, Gary S., Dennis W. Carlton & Hal S. Sider. 2010. "Net Neutrality and Consumer Welfare." *Journal of Competition Law & Economics*, 6(3): 497-519.

Bell Aliant Regional Communications. 2009. Comments to Canadian Radio-Television and Telecommunications Commission. File No. 8646-C12-200815400. February 23. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1029804.zip

Bell, Gregory. 2003. "Failure to Thrive: QoS and the Culture of Operational Networking." *ACM SIGCOMM Workshop on Revisiting IP QoS: What Have We Learned, Why Do We Care?*: 115-120.

Ben-Shahar, Omri & Carl E. Schneider. 2011. "The Failure of Mandated Disclosure." *University of Pennsylvania Law Review*, 159(3): 647-749.

Benkler, Yochai. 2000. "From Consumers to Users: Shifting the Deeper Structures of Regulation Towards Sustainable Commons and User Access." *Federal Communications Law Journal*, 52(3): 561-579.

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

Blake, Steven, David L. Black, Mark A. Carlson, Elwyn Davies, Zheng Wang & Walter Weiss. 1998. "An Architecture for Differentiated Services." Request for Comments 2475. IETF.

Blumenthal, Marjory S. & David D. Clark. 2001. "Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World." *ACM Transactions on Internet Technology*, 1(1): 70-109.

Bode, Karl. 2006. "Which ISPs Throttle BitTorrent Traffic?" *Broadband DSLReports.com.* December 22. http://www.dslreports.com/shownews/80468

Body of European Regulators for Electronic Communications. 2012a. *BEREC Findings on Traffic Management Practices in Europe.* Body of European Regulators for Electronic Communications. BoR (12) 30. http://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/45-berec-findings-on-traffic-management-pra_0.pdf

Body of European Regulators for Electronic Communications. 2012b. *A View of Traffic Management and Other Practices Resulting in Restrictions to the Open Internet in Europe.* Body of European Regulators for Electronic Communications. BoR (12) 30. http://erg.eu.int/doc/consult/bor12_30.pdf

Borthwick, John. 2010. "Neutrality or Bust." *Techcrunch.* December 19. http://techcrunch.com/2010/12/19/neutrality/

Botanical Society of America. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 10. http://apps.fcc.gov/ecfs/document/view?id=7020922879

Brehm, Jack 1956. "Post-Decision Changes in Desirability of Alternatives." *Journal of Abnormal and Social Psychology*, 52: 384-389.

Bricklin, Dan. 2003. "Why We Don't Need QOS: Trains, Cars, and Internet Quality of Service." August 2. http://www.bricklin.com/qos.htm

Briscoe, Bob, Toby Moncaster & Louise Burness. 2007. "Problem Statement: We Don't Have To Do Fairness Ourselves." Internet Draft IEFT.

Briscoe, Bob, Richard Woundy & Alissa Cooper. 2012. "Congestion Exposure (ConEx) Concepts and Use Cases." Request for Comments 6789. IETF.

Brito, Jerry, Martin E. Cave, Robert W. Crandall, Larry F. Darby, Everett Ehrlich, Jeffrey A. Eisenach, et al. 2010. Net Neutrality Regulation: The Economic Evidence. Ex Parte Declaration to Federal Communications Commission. GN Dkt. No. 09-191. http://apps.fcc.gov/ecfs/document/view?id=7020408754

Burgstahler, L. , K. Dolzer, C. Hauser, J. Jähnert, S. Junghans, C. Macián, et al. 2003. "Beyond Technology: the Missing Pieces for QoS Success." *ACM SIGCOMM Workshop on Revisiting IP QoS: What Have We Learned, Why Do We Care?*: 121-130.

Burnham, Brad. 2010a. "Internet Access Should Be Application-Agnostic." *The Huffington Post.* December 20. http://www.huffingtonpost.com/brad-burnham/internet-access-should-be_b_799028.html

Burnham, Brad. 2010b. "A Threat to Startups." *New York Times.* August 10. http://www.nytimes.com/roomfordebate/2010/8/9/who-gets-priority-on-the-web/a-threat-to-startups

Burnham, Thomas A., Judy K. Frels & Vijay Mahajan. 2003. "Consumer Switching Costs: A Typology, Antecedents, and Consequences." *Journal of the Academy of Marketing Science*, 31(2): 109-126.

CACM Staff. 2012. "BufferBloat: What's Wrong With the Internet?" *Communications of the ACM*, 55(2): 40-47.

Calo, Ryan. 2012. "Against Notice Skepticism in Privacy (and Elsewhere)." *Notre Dame Law Review*, 87(3): 1027-1072.

Canadian Film and Television Production Association. 2009. Comments to Canadian Radio-Television and Telecommunications Commission. File No. 8646-C12-200815400. February 23. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1030120.PDF

Canadian Radio-Television and Telecommunications Commission. 2009a. "Canadian Association of Internet Providers et al. and Vaxination Informatique: Application to Review and Vary Certain Determinations in Telecom Decision 2008-108 Related to Bell Canada's Internet Traffic Management Practices. Telecom Decision CRTC 2009-677". CRTC 2009-677. http://crtc.gc.ca/eng/archive/2009/2009-677.htm

Canadian Radio-Television and Telecommunications Commission. 2009b. "Oral Public Hearing Before the CRTC on Review of the Internet Traffic Management Practices of Internet Service Providers, Telecom Public Notice CRTC 2008-19, on July 9, 2009. Transcript of Proceeding, Volume 4.". CRTC 2008-19. http://www.crtc.gc.ca/eng/transcripts/2009/tt0709.htm

Canadian Radio-Television and Telecommunications Commission. 2009c. "Oral Public Hearing Before the CRTC on Review of the Internet Traffic Management Practices of Internet Service Providers, Telecom Public Notice CRTC 2008-19, on July 10, 2009. Transcript of Proceeding, Volume 5.". CRTC 2008-19. http://www.crtc.gc.ca/eng/transcripts/2009/tt0710.htm

Canadian Radio-Television and Telecommunications Commission. 2009d. "Oral Public Hearing Before the CRTC on Review of the Internet Traffic Management Practices of Internet Service Providers, Telecom Public Notice CRTC 2008-19, on July 13, 2009. Transcript of Proceeding, Volume 6.". CRTC 2008-19. http://www.crtc.gc.ca/eng/transcripts/2009/tt0713.htm

Canadian Radio-Television and Telecommunications Commission. 2009e. "Oral Public Hearing Before the CRTC on Review of the Internet Traffic Management Practices of Internet Service Providers, Telecom Public Notice CRTC 2008-19, on July 14, 2009. Transcript of Proceeding, Volume 7.". CRTC 2008-19. http://www.crtc.gc.ca/eng/transcripts/2009/tt0714.htm

Canadian Radio-Television and Telecommunications Commission. 2009f. "Review of the Internet Traffic Management Practices of Internet Service Providers. Telecom Regulatory Policy". CRTC 2011-657. http://www.crtc.gc.ca/eng/archive/2009/2009-657.htm

Carlton, Dennis W. & Jeffrey M. Perloff. 2005. *Modern Industrial Organization.* 4th ed. Boston, MA: Pearson.

Carpenter, Brian, Randall Atkinson & Hannu Flinck. 2010. "Renumbering Still Needs Work." Request for Comments 5887. IETF.

Carpenter, Brian E. & Kathleen Nichols. 2002. "Differentiated Services in the Internet." *Proceedings of the IEEE*, 90(9): 1479-1494.

Carpenter, Brian & Y. Rekhter. 1996. "Renumbering Needs Work." Request for Comments 1900. IETF.

Cate, Fred H. 2006. "The Failure of Fair Information Practice Principles." In *Consumer Protection in the Age of the 'Information Economy',* ed. Jane K. Winn, 341-378. Hampshire, UK: Ashgate Publishing.

Cave, Martin, Richard Collins, Nico van Eijk, Pierre Larouche, Luigi Propsperetti, Alexandre de Streel, et al. 2009. *Statement by European Academics on the Inappropriateness of Imposing Increased Internet Regulation in the EU.* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1329926

Cave, Martin & Pietro Crocioni. 2007. "Does Europe Need Network Neutrality Rules?" *International Journal of Communication*, 1: 669-679.

Cave, Martin & Pietro Crocioni. 2011. "Net Neutrality in Europe." *Communications & Convergence Review*, 3(1): 57-70.

Cellan-Jones, Rory. 2009. "iPlayer: BBC v BT." *BBC News dot.life Blog.* June 2. http://www.bbc.co.uk/blogs/technology/2009/06/iplayerbbc_v_bt.html

Center for Democracy & Technology. 2010. Comments to Federal Communications Commission. GN Docket No. 09-191. January 14. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020378292

Center for Media Justice, Consumers Union, Media Access Project, New America Foundation & Public Knowledge. 2010. Comments of Public Interest Commenters to Federal Communications Commission. GN Docket No. 09-191. January 14. http://apps.fcc.gov/ecfs/document/view?id=7020378818

Cerf, Vinton G. 2006. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: Network Neutrality.* commerce.senate.gov/pdf/cerf-020706.pdf

Challenge.gov. 2011. "FCC Open Internet Apps Challenge." August 5. http://openinternetapps.challenge.gov/

Chen, Brian X. 2012. "AT&T Chief Regrets Offering Unlimited Data for iPhone." *New York Times.* May 4. http://bits.blogs.nytimes.com/2012/05/04/att-randall-stephenson/

Chu, Jerry, Nandita Dukkipati, Yuchung Cheng & Matt Mathis. 2013. "Increasing TCP's Initial Window." Internet Draft IETF.

Cisco Systems. 2005a. *Network-Based Application Recognition and Distributed Network-Based Application Recognition Feature Guide (Cisco IOS Release 12.4(4) T).* Cisco. http://www.cisco.com/en/US/docs/ios/12_1/12_1e11/feature/guide/dtnbarad.pdf

Cisco Systems. 2005b. "Quality of Service Design Overview." In *Enterprise QoS Solution Reference Network Design Guide,*Vol. 3.3, 1/1-1/34. San Jose, CA: Cisco Systems.

Cisco Systems. 2012. *NBAR2 Protocol Library (Last Update: January 2012).* Product Bulletin PB627831. http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6558/ps6616/product_bulletin_c25-627831.pdf

Cisco Systems, Inc. 2010. Comments to Federal Communications Commission. GN Docket No. 09-191. January 14. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020374147

Clark, David D., John Wroclawski, Karen R. Sollins & Robert Braden. 2005. "Tussle in Cyberspace: Defining Tomorrow's Internet." *IEEE/ACM Transactions on Networking*, 13(3): 462- 475.

Clark, David, William Lehr & Steven Bauer. 2010. Comments to Federal Communications Commission. GN Dkt. No. 09-191. January 14. http://apps.fcc.gov/ecfs/document/view?id=7020373725

Clyburn, Mignon L. 2010. "Preserving the Open Internet. Report and Order. Statement of Commissioner Mignon L. Clyburn, Approving in Part, Concurring in Part". FCC 10-201. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A5.pdf

Cohen, David. 2008. *Written Testimony before the Federal Communications Commission at its First En Banc Hearing on Broadband Management Practices.* Federal Communications Commission. http://www.fcc.gov/broadband_network_management/022508/cohen.pdf

Comcast Corporation. 2008a. *Comcast Corporation Description of Current Network Management Practices.* Attachment A to Comcast Corporation's Filing In the Matter of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for

Secretly Degrading Peer-to-Peer Applications submitted September 19, 2008. WC Dkt. No. 07-52. http://apps.fcc.gov/ecfs/document/view?id=6520172537

Comcast Corporation. 2008b. *Comcast Corporation Description of Planned Network Management Practices to be Deployed Following the Termination of Current Practices.* Attachment B to Comcast Corporation's Filing In the Matter of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications submitted September 19, 2008. WC Dkt. No. 07-52. http://apps.fcc.gov/ecfs/document/view?id=6520172538 and http://apps.fcc.gov/ecfs/document/view?id=6520172538

Comcast Corporation. 2012a. "About Excessive Use of Data." *Comcast Customer Central.* May 17. http://customer.comcast.com/help-and-support/internet/common-questions-excessive-use

Comcast Corporation. 2012b. "XFINITY TV. FAQs: Xbox 360." http://xbox.comcast.net/faqs.html

Consumer's Union. 2007. "The Next Big Thing in Broadband: Early Termination Penalties." *HearUsNow.org Blog.* April 9. http://www.consumersunion.org/blogs/hun/2007/04/the_next_big_thing_in_broadban.html.

Cooper, Alissa. 2011. "Doing the DPI Dance." In *Privacy in America: Interdisciplinary Perspectives,* eds. William Aspray & Philip Doty, 139-165. Lanham, MD: Scarecrow Press, Inc.

Cooper, Alissa. 2013a. "How Competition Drives Discrimination: An Analysis of Broadband Traffic Management in the UK." Paper presented at 41st Research Conference on Communication, Information and Internet Policy (TPRC 41). Arlington, Virginia, USA.

Cooper, Alissa. 2013b. "How Regulation and Competition Influence Discrimination in Broadband Traffic Management: A Comparative Study of Net Neutrality in the United States and the United Kingdom." DPhil Thesis. Oxford University, Oxford, UK.

Copps, Michael J. 2010. "Preserving the Open Internet. Report and Order. Concurring Statement of Commissioner Michael J. Copps". FCC 10-201. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A3.pdf

Corbató, Steven C. & Ben Teitelbaum. 2006. *Internet2 and Quality of Service: Research, Experience, and Conclusions.* Educause. http://www.educause.edu/ir/library/pdf/CSD4577.pdf

Covad Communications Company. 2010. Comments to Federal Communications Commission. GN Docket No. 09-191. January 14. http://apps.fcc.gov/ecfs/document/view?id=7020373797

Cox Communications, Inc. 2010. Comments to Federal Communications Commission. GN Docket No. 09-191. January 14. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020378714

Cox Communications Inc. 2009. "Congestion Management FAQs." *Cox.* February 13. http://www.cox.com/policy/congestionmanagement/

Craswell, Richard. 1982. "Tying Requirements in Competitive Markets: The Consumer Protection Issues." *Boston University Law Review*, 62(3): 661-700.

Crawford, Susan. 2007. "The Internet and the Project of Communications Law." *UCLA Law Review*, 55(2): 359-407.

Crawford, Susan. 2009. "Transporting Communications." *Boston University Law Review*, 89(3): 871-938.

Crawford, Susan. 2010. "The Looming Cable Monopoly." *Yale Law and Policy Review Inter Alia*, 29: 34-40.

Crawford, Susan. 2011. "The Communications Crisis in America." *Harvard Law & Policy Review*, 5(2): 245-263.

Croley, Steven P. 1998. "Theories of Regulation: Incorporating the Administrative Process." *Columbia Law Review*, 98(1): 1-168.

Cullen, Joseph & Oleksandr Shcherbakov. 2010. *Measuring Consumer Switching Costs in the Wireless Industry.* Attachment 1 to the Open Internet Coalition's Reply Comments In the Matter of Preserving the Open Internet submitted April 26, 2010. GN Dkt. No. 09-191. http://apps.fcc.gov/ecfs/document/view?id=7020438577

Davie, Bruce. 2003. "Deployment Experience with Differentiated Services." *Proceedings of the ACM SIGCOMM Workshop on Revisiting IP QoS: What have we learned, why do we care? (RIPQoS '03)*: 131-136.

Davies, Neil, Judy Holyer & Peter Thompson. 1999. "An Operational Model to Control Loss and Delay of Traffic at a Network Switch." Paper presented at Third IFIP Workshop on Traffic Management and Design of ATM Networks.

DellaVigna, Stefano & Ulrike Malmendier. 2004. "Contract Design and Self-Control Theory and Evidence." *The Quarterly Journal of Economics*, CXIX(2): 353-402.

Dischinger, Marcel, Alan Mislove, Andreas Haeberlen & Krishna P. Gummadi. 2008. "Detecting BitTorrent Blocking." *Internet Measurement Conference 2008*: 3-8.

Distributel Communications Limited. 2009. Comments to Canadian Radio-Television and Telecommunications Commission. File No. 8662-V42-200907826. June 22. http://crtc.gc.ca/public/partvii/2009/8662/p8_200907727/1231235.pdf

Documentary Organization of Canada. 2009. Comments to Canadian Radio-Television and Telecommunications Commission. File No. 8646-C12-200815400. February 23. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1030141.PDF

Draft Bill To Amend Title I of the Communications Act of 1934 to Provide for Internet Openness, and for Other Purposes. 2010. Attachment to Henry Waxman's Ex Parte Letter to Federal Communications Comission submitted December 1, 2010. GN Docket No. 09-191. http://apps.fcc.gov/ecfs/document/view?id=7020922223

Dugan, Andrew. 2012. "An IP Engineer and Consumer View of Xfinity Traffic Prioritization." *Beyond Bandwith.* May 17. http://blog.level3.com/2012/05/17/an-ip-engineer-and-consumer-view-of-xfinity-traffic-prioritization/

Dunbar, John. 2007. "Pulling Plug on Net Service Not Easy." *USA Today.* April 9. http://www.usatoday.com/tech/techinvestor/industry/2007-04-09-net-fees_N.htm?POE=TECISVA.

Eckersley, Peter, Fred von Lohmann & Seth Schoen. 2007. *Packet Forgery by ISPs: A Report on the Comcast Affair.* Electronic Frontier Foundation. http://www.eff.org/files/eff_comcast_report.pdf

Economides, Nicholas. 2010. *Why Imposing New Tolls on Third-Party Content and Applications Threatens Innovation and will not Improve Broadband Providers' Investment.* Appendix A of Google's Comments to Federal Communications Commission submitted January 14, 2010 GN Docket No. 09-101. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020378726

Edwards, Matthew A. 2005. "Empirical and Behavioral Critiques of Mandatory Disclosure: Socio-Economics and the Quest for Truth in Lending." *Cornell Journal of Law and Public Policy*, 14(2): 199-249.

Eggerton, John. 2010. "Free Press Pushed Hard Against Waxman's Net Neutrality Bill." *Multichannel News.* November 30. http://www.multichannel.com/article/457836-Free_Press_Pushed_Hard_Against_Waxman_s_Net_Neutrality_Bill.php

Electronic Frontier Foundation. 2008. "Test Your ISP." *Electronic Frontier Foundation.* October 26. https://www.eff.org/testyourisp

Engelhart, Ken. 2007. "Cyber Traffic Moving Smoothly." *The Star.* April 20. http://www.thestar.com/article/205204

Eskridge, Jr., William N. 1988. "Politics without Romance: Implications of Public Choice Theory for Statutory Interpretation." *Virginia Law Review*, 74(2): 275-338.

European Commission. 2007. "Commission Staff Working Document. Impact Assessment". SEC(2007) 1472. http://ec.europa.eu/information_society/policy/ecomm/doc/library/proposals/1472/comm_pdf_sec_2007_1472_1_en_documentdetravail.pdf.pdf

European Commission. 2011. "The Open Internet and Net Neutrality in Europe". COM(2011) 222 final. http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2011:0222:FIN:EN:PDF

Farber, David, Michael L. Katz, Gerald Faulhaber & Christopher S. Yoo. 2007. "Hold Off on Network Neutrality." *Washington Post.* January 19. http://www.washingtonpost.com/wp-dyn/content/article/2007/01/18/AR2007011801508.html

Farrell, Joseph. 2006. "Open Access Arguments: Why Confidence Is Misplaced." In *Net Neutrality or Net Neutering: Should Broadband Internet Services Be Regulated?,* eds. Thomas M. Lenard & Randolph J. May, 195-214. New York, NY: Springer.

Farrell, Joseph & Paul Klemperer. 2007. "Coordination and Lock-in : Competition and Switching Costs and Network Effects." In *Handbook of Industrial Organization,* eds. Mark Armstrong & Robert Porter, Vol. 3, 1967-2072. Amsterdam, The Netherlands: Elsevier.

Faulhaber, Gerald. 2009. "A National Broadband Plan for Our Future: A Customer-Centric Framework." *International Journal of Communication*, 3: 742-779.

Faulhaber, Gerald R. . 2011. "Economics of Net Neutrality: A Review." *Communications & Convergence Review*, 3(1): 7-25.

Faulhaber, Gerald R. & David J. Farber. 2010a. "Innovation in the Wireless Ecosystem: A Customer-Centric Framework." *International Journal of Communication*, 4: 73-112.

Faulhaber, Gerald R. & David J. Farber. 2010b. "The Open Internet: A Customer-Centric Framework." *International Journal of Communication*, 4: 302-342.

Feamster, Nick. 2010. "Outsourcing Home Network Security " *2010 ACM Sigcomm Workshop on Home Networks, HomeNets'10*: 37-42.

Federal Communications Commission. 2005. "Appropriate Framework for Broadband Access to the Internet over Wireline Facilities. Policy Statement". *F.C.C.R.,* 20: 14986-14988

Federal Communications Commission. 2007. "Public Notice Asking for Comments and Reply Comments regarding Petition for Rulemaking in the Matter of E-mail Address Portability". Report No. 2832. http://apps.fcc.gov/ecfs/document/view?id=6519739081

Federal Communications Commission. 2008. "Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications. Memorandum Opinion and Order". FCC 08-183.

Federal Communications Commission. 2009a. "Letter to Comcast In the Matters of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications". WC 07-52. http://fjallfoss.fcc.gov/ecfs/document/view?id=6520213534

Federal Communications Commission. 2009b. "Preserving the Open Internet. Notice of Proposed Rulemaking". FCC 09-93. hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-09-93A1.pdf

Federal Communications Commission. 2010a. *Broadband Decisions: What Drives Consumers to Switch – or Stick with – Their Broadband Internet Provider.* FCC Working Paper. http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-303264A1.pdf

Federal Communications Commission. 2010b. "Connecting America: The National Broadband Plan". FCC 09-31. http://download.broadband.gov/plan/national-broadband-plan.pdf

Federal Communications Commission. 2010c. "Preserving the Open Internet. Report and Order". FCC 10-201. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf

Federal Communications Commission. 2011a. "47 CFR Parts 0 and 8. GN Docket No. 09–191; WC Docket No. 07–52; FCC 10–201. Preserving the Open Internet. Final Rule". *Federal Register,* 76: 59192-59235. http://www.gpo.gov/fdsys/pkg/FR-2011-09-23/pdf/2011-24259.pdf

Federal Communications Commission. 2011b. "Notice to the U.S. Judicial Panel on Multicircuit Litigation. Notice of Multicircuit Petitions for Review". http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-310240A1.pdf

Federal Communications Commission. 2014. "Protecting and Promoting the Open Internet. Notice of Proposed Rulemaking". FCC 14-61. http://transition.fcc.gov/Daily_Releases/Daily_Business/2014/db0515/FCC-14-61A1.pdf

Federal Trade Commission. 2003. "Federal Trade Commission Comments to be Filed in Dockets OST-97-2881, OST-97-3014 and OST-98-4775." *Federal Trade Commission.* June 6. http://ftc.gov/os/2003/06/dotcomment.htm

Feld, Harold. 2010a. "Genachowski Hits The Legal Reset Button — "Title II Lite"." *Wetmachine - Tales of the Sausage Factory.* May 7. http://tales-of-the-sausage-factory.wetmachine.com/content/genachowski-hits-the-legal-reset-button-title-ii-lite

Feld, Harold. 2010b. "Genachowski's Fast Fading Star — And How He Can Still Salvage His Term As Chairman." *Wetmachine - Tales of the Sausage Factory.* August 3. http://tales-of-the-sausage-factory.wetmachine.com/genachowskis-fast-fading-star-and-how-he-can-still-salvage-his-term-as-chairman/

Feld, Harold. 2010c. "Sorry AT&T, Title II Would Not Require The FCC To Allow Paid Prioritization." *Wetmachine - Tales of the Sausage Factory.* October 8. http://www.wetmachine.com/content/sorry-att-title-ii-would-not-require-the-fcc-to-allow-paid-prioritization

Frederick, Shane , George Loewenstein & Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40(2): 351-401.

Free Press. 2009a. Comments to Federal Communications Commission. GN Docket No. 09-137. September 4. http://apps.fcc.gov/ecfs/document/view?id=7020037663

Free Press. 2009b. Reply Comments to Federal Communications Commission. GN Dkt. No. 09-51. July 21. http://apps.fcc.gov/ecfs/document/view?id=7019917712

Free Press. 2010. Reply Comments to Federal Communications Commission. GN Docket No. 09-191. April 26. http://apps.fcc.gov/ecfs/document/view?id=7020437465

Free Press & Public Knowledge. 2007. Formal Complaint to Federal Communications Commission. November 1

Frischmann, Brett. 2012. *Infrastructure: The Social Value of Shared Resources.* New York, NY: Oxford University Press.

Frischmann, Brett M. 2005. "An Economic Theory of Infrastructure and Commons Management." *Minnesota Law Review*, 89 (April): 917-1030.

Frischmann, Brett M. 2010 Ex Parte Comments to Federal Communications Commission. GN Dkt. No. 09-191. April 12. http://fjallfoss.fcc.gov/ecfs2/document/view.action?id=7020408747

Frischmann, Brett M. & Barbara van Schewick. 2007. "Network Neutrality and the Economics of an Information Superhighway: A Reply to Professor Yoo." *Jurimetrics Journal*, 47(4): 383–428.

Geist, Michael. 2007. "ISP Must come Clean on 'Traffic Shaping'." *The Star.* April 16. http://www.thestar.com/sciencetech/article/203408

Geist, Michael. 2011. "CRTC's Net Neutrality Rules in Action: Bell To Drop P2P Traffic Shaping." *Michael Geist Blog.* December 20. http://www.michaelgeist.ca/content/view/6209/125/

Genachowski, Julius. 2010. *Remarks on Preserving Internet Freedom and Openness.* http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-303136A1.pdf

Gettys, Jim. 2011. "IW10 Considered Harmful." Internet Draft IETF.

Gettys, Jim & Kathleen Nichols. 2012. "Bufferbloat: Dark Buffers in the Internet." *Communications of the ACM*, 55(1): 57-65.

Google Inc. 2010. Reply Comments to Federal Communications Commission. GN Docket No. 09-191. April 26. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020438889

Google, Inc. & Inc. Verizon. 2010. "Google-Verizon Legislative Framework Proposal." August 9. http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/googleblogs/pdfs/verizon_google_legislative_framework_proposal_081010.pdf

Gross, Grant. 2010. "Net Neutrality Plan Has the Votes at FCC." *Computerworld.* December 20. http://www.computerworld.com/s/article/9201818/Net_neutrality_plan_has_the_votes_at_FCC?taxonomyId=70&pageNumber=1

Hahn, Robert, Robert Litan & Hal J. Singer. 2010. "Addressing the Next Wave of Internet Regulation: Toward a Workable Principle for Nondiscrimination." *Regulation & Governance*, 4(3): 365-382.

Hahn, Robert W. & Scott J. Wallsten. 2006. "The Economics of Net Neutrality." *The Economist's Voice*, 3(6): Article 8.

Hancock, Robert, Georgios Karagiannis, John Loughney & Sven Van den Bosch. 2005. "Next Steps in Signaling (NSIS): Framework." Request for Comments 4080. IETF.

Hart, Kim. 2010a. "Net Neutrality Bill Stillborn." *Politico.* September 29. http://www.politico.com/news/stories/0910/42919.html

Hart, Kim. 2010b. "Sources: FCC Chief to Move on Net Neutrality." *Politico.* November 18. http://www.politico.com/news/stories/1110/45371.html

Hatch, David. 2010. "FCC's Copps Isn't Committed On Net Neutrality." *National Journal Daily.* December 9. http://www.nationaljournal.com/member/daily/fcc-s-copps-isn-t-committed-on-net-neutrality-20101209

Hausman, Jerry A., Gregory J. Sidak & Hal J. Singer. 2001. "Residential Demand for Broadband Telecommunications and Consumer Access to Unaffiliated Internet Content Providers." *Yale Journal on Regulation*, 18(1): 129-173.

Hazlett, Thomas W. & Joshua D. Wright. 2011. *The Law and Economics of Network Neutrality (Version September 12, 2011; last revised November 2, 2011).* George Mason University George Mason Law & Economics Research Paper 11-36

Hellweg, Eric. 2003. The Kazaa Conundrum: The RIAA's Lawsuits Have Split the Broadband ISP Industry and Will Likely Split Investors as Well. *CNN Money.com.* September 10. http://money.cnn.com/2003/09/10/technology/techinvestor/hellweg/index.htm.

Hemphill, C. Scott. 2008. "Network Neutrality and the False Promise of Zero-Price Regulation." *Yale Journal on Regulation*, 25(2): 135-180.

Herman, Bill D. 2006. "Opening Bottlenecks: On Behalf of Mandated Network Neutrality." *Federal Communications Law Journal*, 59(1): 103-156.

Hesseldahl, Arik. 2010. "Breaking: FCC's Copps Voting "Yes" on Net Neutrality Plan." *All Things D.* December 20. http://allthingsd.com/20101220/breaking-fcc-commissioner-copps-says-hell-vote-yes/

Higginbotham, Stacey. 2012. "He Said, She Said: Is Comcast Prioritizing Traffic or Not?" *GigaOm.* May 15. http://gigaom.com/broadband/he-said-she-said-is-comcast-prioritizing-traffic-or-not/

Hogendorn, Christiaan. 2012. "Spillovers and Network Neutrality." In *Regulation and the Performance of Communication and Information Networks* eds. Gerald R. Faulhaber, Gary Madden & Jeffrey Petchey, 191-208. Cheltenham Glos, UK: Edward Elgar Publishing.

Hovenkamp, Herbert. 2011a. *Antitrust Law: An Analysis of Antitrust Principles and Their Application* 3rd ed. Vol. XI. New York, NY: Aspen Publishers.

Hovenkamp, Herbert. 2011b. *Federal Antitrust Policy: The Law of Competition and Its Practice.* 4th ed. St. Paul, MN: West Publishing Co. .

Hultquist, Hank. 2010. "Harold Feld is Right (About Some Things)." *AT&T Public Policy Blog.* October 13. http://attpublicpolicy.com/government-policy/harold-feld-is-right-about-some-things/

Huston, Geoff. 2012. "The QoS Emperor's Wardrobe." *The ISP Column.* June. http://www.potaroo.net/ispcol/2012-06/noqos.html

IEEE-USA. 2010. *Position Statement Network Traffic Management. Adopted by the IEEE-USA Board of Directors (19 Nov. 2010).* http://ieeeusa.org/policy/positions/NetworkTrafficManagementNov10.pdf

IEEE-USA Committee on Communications Policy. 2010. *Network Traffic Management and the Evolving Internet. White Paper.* IEEE-USA. www.ieeeusa.org/policy/whitepapers/IEEEUSAWP-NTM2010.pdf

Independent Film & Television Alliance. 2009. Comments to Canadian Radio-Television and Telecommunications Commission. File No. 8646-C12-200815400. February 23. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1032007.DOC

International Telecommunication Union. 2003. *ITU-T Recommendation G.114: One-way transmission time.* International Telecommunication Union. G.114. http://www.itu.int/rec/dologin_pub.asp?lang=e&id=T-REC-G.114-200305-I!!PDF-E&type=items

Internet2. Undated. "QBone Scavenger Service (QBSS)." *QBone Website.* http://qbone.internet2.edu/qbss/

Internet Engineering Task Force. 2010. "Congestion Exposure (conex) Working Group Charter." https://datatracker.ietf.org/wg/conex/charter/

Internet Freedom and Nondiscrimination Act of 2006. H.R. 5417. 109th Congress, 2nd Session. 2006. http://www.govtrack.us/congress/bills/109/hr5417/text

Internet Freedom Preservation Act. S. 215. 110th Congress, 1st Session. 2007. http://www.govtrack.us/congress/bills/110/s215/text

Internet Non-Discrimination Act of 2006. S. 2360. 109th Congress, 2nd Session. 2006. https://www.govtrack.us/congress/bills/109/s2360/text

Jacobson, Van, Bob Braden & Dave Borman. 1992. "TCP Extensions for High Performance." RFC 1323. IETF.

Jacquet, Arnaud, Bob Briscoe & Toby Moncaster. 2008. "Policing Freedom to Use the Internet Resource Pool " Paper presented at ACM CoNEXT 2008. Madrid, Spain.

Järvinen, Ilpo, Binoy Chemmagate, Laila Daniel, Aaron Yi Ding, Markku Kojo & Markus I Isomäki. 2012. "Impact of TCP on Interactive Real-Time Communication." Paper presented at IAB / IRTF Workshop on Congestion Control for Interactive Real-Time Communication. Vancouver, Canada.

Jerome, Sara. 2010a. "Sources: OIC Not Supporting Waxman Net-Neutrality Bill." *The Hill.* September 28. http://thehill.com/blogs/hillicon-valley/technology/121493-sources-oic-will-not-support-waxman-net-neutrality-bill

Jerome, Sara. 2010b. "Waxman May Move Without Full Consensus on Net-Neutrality Bill." *The Hill.* September 23. http://thehill.com/blogs/hillicon-valley/technology/120571-waxman-may-move-without-full-consensus-on-net-neutrality-bill

Jordan, Scott & Arijit Ghosh. 2010. "A Framework for Classification of Traffic Management Practices as Reasonable or Unreasonable." *ACM Transactions on Internet Technology*, 10(3): 12:11-12:23.

Kahneman, Daniel & Amos Tversky. 1984. "Choices, Values and Frames." *American Psychologist*, 39(4): 341-350.

Kang, Cecilia. 2009. "AT&T, Carriers Fund Democratic Reps Against Net Neutrality." *Washington Post.* October 19. http://voices.washingtonpost.com/posttech/2009/10/all_but_two_of_the.html

Kang, Cecilia. 2010a. "FCC Draws Fire Over Talks with Internet, Telecom Giants on 'Net Neutrality'." *Washington Post.* Augustt 5. http://www.washingtonpost.com/wp-dyn/content/article/2010/08/04/AR2010080406910_pf.html

Kang, Cecilia. 2010b. "FCC Ends Talks for Deal on Net Neutrality." *Washington Post.* August 6. http://www.washingtonpost.com/wp-dyn/content/article/2010/08/05/AR2010080502423_pf.html

Kang, Cecilia. 2012. "AT&T Faces Complaint Over iPhone FaceTime Blocking." *Washington Post.* September 19. http://www.washingtonpost.com/blogs/post-tech/post/atandt-faces-complaint-over-iphone-facetime-blocking/2012/09/18/799c8650-0183-11e2-b257-e1c2b3548a4a_blog.html

Kang, Jerry. 2007. "Race.Net Neutrality." *Journal on Telecommunications and High Technology Law*, 6(1): 1-22.

Kanter, James & Mark Scott. 2014. "European Panel Adopts 'Net Neutrality' and Mobile Roaming Rules." *New York Times.* March 18. http://www.nytimes.com/2014/03/19/business/international/eu-panel-adopts-net-neutrality-and-mobile-roaming-rules.html

Kaplow, Louis. 1992. "Rules Versus Standards: An Economic Analysis." *Duke Law Journal*, 42(3): 557-629.

Kaplow, Louis. 1999. "General Characteristics of Rules." In *Encyclopedia of Law and Economics,* eds. Boudewijn Bouckaert & Gerrit De Geest, 502-528. Northampton, MA: Edward Elgar Publishing.

Karr, Timothy. 2010. "Netroots to Obama FCC: Inaction Is Not an Option." *The Huffington Post.* May 4. http://www.huffingtonpost.com/timothy-karr/netroots-to-obama-fcc-ina_b_562564.html

Kopf, Jared, Jeff Bezos, Ashwin Navin, James F. Geiger, Craig Newmark, Jay Adelson, et al. 2010. *Letter to Chairman Genachowski.* Appendix A to the Open Internet Coalition's Comments In the Matter of Preserving the Open Internet submitted January 14, 2010. GN Dkt. No. 09-191. http://apps.fcc.gov/ecfs/document/view?id=7020377928

Krafft, Jackie & Evens Salies. 2008. "The Diffusion of ADSL and Costs of Switching Internet Providers in the Broadband Industry: Evidence from the French Case." *Research Policy*, 37(4): 706-719.

Krafft, Jackie & Evens Salies. 2009. "Why and How Should New Industries with High Consumer Switching Costs be Regulated? The Case of Broadband Internet in France." In *Regulation, Deregulation, Reregulation - Institutional Perspectives,* eds. Claude Menard & Michel Ghertman, 327-350. Northampton, MA: Edward Elgar Publishing.

Kroes, Neelie. 2012. "Next Steps on Net Neutrality – Making Sure you get Champagne Service if That's What You're Paying For." *European Commission.* May 29. http://blogs.ec.europa.eu/neelie-kroes/netneutrality/

Kuhns, Jeff C. 2006. *Testimony Before the United States Senate Committee on the Judiciary at its Hearing on Reconsidering Our Communications Laws: Ensuring Competition and Innovation.* http://www.judiciary.senate.gov/hearings/testimony.cfm?id=e655f9e2809e5476862f735da116309f&wit_id=e655f9e2809e5476862f735da116309f-2-7

Kurose, James F. & Keith W. Ross. 2010. *Computer Networking: A Top-Down Approach.* 5th ed. Boston, MA: Pearson/Addison Wesley.

LambdaJive. 2009. "iSinglePayer iPhone App Censored by Apple." *LambdaJive.* September 26. http://lambdajive.wordpress.com/2009/09/26/isinglepayer-iphone-app-censored-by-apple/

Lasar, Matthew. 2009. "AT&T and Astroturf: Is "Following the Money" Enough?" *Ars Technica.* October 26. http://arstechnica.com/tech-policy/2009/10/the-anti-net-neutrality-movement-is-it-just-about-att-money/

Lasar, Matthew. 2010. "Hey, Capitol Hill: Who's your Daddy? AT&T." *Ars Technica.* September 14. http://arstechnica.com/tech-policy/2010/09/hey-capitol-hill-whose-your-daddy-att/

Latin, Howard. 1994. "'Good' Warnings, Bad Products, and Cognitive Limitations." *UCLA Law Review*, 41(5): 1193-1295.

Lee, Robin S. & Tim Wu. 2009. "Subsidizing Creativity through Network Design: Zero-Pricing and Net Neutrality." *Journal of Economic Perspectives*, 23(3): 61-76.

Lehr, William H., Sharon E. Gillett, Jon M. Peha & Marvin A. Sirbu. 2007. "Scenarios for the Network Neutrality Arms Race." *International Journal of Communication*, 1: 607-643.

Leichtman Research Group. 2010. "Over 800,000 Add Broadband in the Third Quarter of 2010." November 15. http://www.leichtmanresearch.com/press/111510release.html

Leichtman Research Group. 2011. "Over 635,000 Add Broadband in the Third Quarter of 2011." November 18. http://www.leichtmanresearch.com/press/111811release.html

Leichtman Research Group. 2014. "Nearly 1.2 Million Add Broadband in the First Quarter of 2014." May 20. http://www.leichtmanresearch.com/press/052014release.html

Lemley, Mark A. & Lawrence Lessig. 1999. Ex Parte to Federal Communications Commission. CS Dkt. No. 99-251. November 10. http://fjallfoss.fcc.gov/ecfs/document/view?id=6009850930

Lennett, Benjamin. 2009. "Dis-empowering Users vs. Maintaining Internet Freedom: Network Management and Quality of Service (QoS)." *The CommLaw Conspectus: Journal of Communications Law and Policy*, 18(1): 97-147.

Lessig, Lawrence. 2001a. *The Future of Ideas: The Fate of the Commons in a Connected World.* New York, NY: Vintage Books.

Lessig, Lawrence. 2001b. "The Internet Under Siege." *Foreign Policy.* November 1. http://www.foreignpolicy.com/articles/2001/11/01/the_internet_under_siege

Lessig, Lawrence. 2002. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: The Government's Role in Promoting the Future of Telecommunications Industry and Broadband Deployment.* http://commerce.senate.gov/pdf/lessig.pdf

Lessig, Lawrence. 2006. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: Network Neutrality.* www.lessig.org/blog/archives/Lessig_Testimony_2.pdf

Lessig, Lawrence. 2008a. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. August 20. http://cyberlaw.stanford.edu/files/blogs/LessigFccComcastLetter.pdf

Lessig, Lawrence. 2008b. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: The Future of the Internet.* http://commerce.senate.gov/public/?a=Files.Serve&File_id=da3dac3c-6a90-4467-96fe-2abbeb78456b

Linux Tracker. 2012. "Linux Tracker." January 17. http://linuxtracker.org/

Litan, Robert E. & Hal J. Singer. 2007. "Unintended Consequences of Net Neutrality Regulation." *Journal on Telecommunications & High Technology Law*, 5(3): 533-572.

Loewenstein, George & Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *The Quarterly Journal of Economics*, 107(2): 573-597.

Lublin, Joann S. & Spencer E. Ante. 2011. "Inside Sprint's Bet on iPhone." *Wall Street Journal.* October 4. http://online.wsj.com/article/SB10001424052970203405504576603053795839250.html

Lyons, Daniel A. 2013. "Net Neutrality and Nondiscrimination Norms in Telecommunications." *Arizona Law Review*, 54(4): 1029-1071.

Magill, M. Elizabeth. 2004. "Agency Choice of Policymaking Form." *The University of Chicago Law Review*, 71(4): 1381-1447.

Manner, Jukka, Georgios Karagiannis & Andrew McDonald. 2010. "NSIS Signaling Layer Protocol (NSLP) for Quality-of-Service Signaling." Request for Comments 5974. IETF.

Marcus, J. Scott. 2008. "Network Neutrality: The Roots of the Debate in the United States." *Intereconomics*, 43(1): 30-37.

Marcus, J. Scott, Pieter Nooren, Jonathan Cave & Kenneth R. Carter. 2011. *Network Neutrality: Challenges and Responses in the EU and in the U.S. Study requested by the European Parliament's Committee on the Internal Market and Consumer Protection.* European Parliament IP/A/IMCO/ST/2011-02. Publication Number PE457.369. http://www.europarl.europa.eu/document/activities/cont/201108/20110825ATT25266/20110825ATT25266EN.pdf

Marcus, J. Scott, Christian Wernick & Kenneth R. Carter. 2008. *Network Neutrality: Implications for Europe.* WIK Diskussionsbeitrag 314. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1522039

Mayer, Lindsay Renick. 2009. "Big Donors Ramp Up to Fight FCC Net Neutrality Decision." *OpenSecrets.org.* October 28. http://www.opensecrets.org/news/2009/10/the-federal-communications-com.html

Mayo, John W., Marius Schwartz, Bruce Owen, Robert Shapiro, Lawrence J. White & Glenn Woroch. 2010. "How to Regulate the Internet Tap." *The New York Times.* April 20, 2010. http://www.nytimes.com/2010/04/21/opinion/21mayo.html

Mennecke, Thomas. 2005. "DSL Broadband Providers Perform Balancing Act." *Slyck News.* November 1. http://www.slyck.com/news.php?story=973.

Mortenson, Gail M. 2007. Petition for Rulemaking to Federal Communications Commission. Dkt. No. RM-11391. July 20. http://fjallfoss.fcc.gov/ecfs/document/view?id=6519560444

National Cable & Telecommunications Association. 2007. Comments to Federal Communications Commission. WC Dkt. No. 07-52. June 15. http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6519529328

Network Neutrality Act of 2006. H.R. 5273. 109th Congress 2nd Session. 2006. https://www.govtrack.us/congress/bills/109/hr5273/text

New York Times. 2010. "The Price of Broadband Politics. Editorial." *New York Times.* June 29. http://www.nytimes.com/2010/06/30/opinion/30wed2.html?_r=1

Nichols, Kathleen. 2012. "Controlling Queue Delay." *ACM Queue*, 10(5).

Nichols, Kathleen, Steven Blake, Fred Baker & David L. Black. 1998. "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers." Request for Comments 2474. IETF.

North American Benthological Society. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 13. http://apps.fcc.gov/ecfs/document/view?id=7020923349

Nuechterlein, Jonathan E. . 2009. "Antitrust Oversight of an Antitrust Dispute: An Institutional Perspective on the Net Neutrality Debate." *Journal on Telecommunications and High Technology Law*, 7(1): 19-65.

NYSERNet. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 13. http://apps.fcc.gov/ecfs/document/view?id=7020923371

Obama for America. 2007. "Barack Obama: Connecting and Empowering All Americans Through Technology and Innovation." November 14. http://my.barackobama.com/page/-/HQpress/111307%20Innovation%20fact%20sheet.pdf

Odlyzko, Andrew. 1998. *The Economics of the Internet: Utility, Utilization, Pricing, and Quality of Service.* Unpublished Manuscript. http://www.dtc.umn.edu/~odlyzko/doc/internet.economics.pdf

Odlyzko, Andrew. 1999. "The Current State and Likely Evolution of the Internet." *IEEE Global Telecommunications Conference 1999 (IEEE GLOBECOM '99)*: 1869-1876.

Odlyzko, Andrew. 2008. "The Delusions of Net Neutrality." Paper presented at 36th Research Conference on Communication, Information and Internet Policy (TPRC 2008).

Ohlhausen, Maureen K. 2012. "The Open Internet: Regulating to Save the Unregulated Internet? Remarks before the MaCCI Law and Economics Conference on The Future of the Internet (October 26, 2012)." http://www.ftc.gov/sites/default/files/documents/public_statements/open-internet-regulating-save-unregulated-internet/121026mannheim_0.pdf

Ohm, Paul. 2009. "The Rise and Fall of Invasive ISP Surveillance." *University of Illinois Law Review*, 2009(5): 1417-1496.

Olson, Mancur. 1980. *The Logic of Collective Action.* Cambridge, MA: Harvard University Press.

Open Internet Coalition. 2010. Comments to Federal Communications Commission. GN Dkt. No. 09-191. January 14. http://apps.fcc.gov/ecfs/document/view?id=7020377928

OpenOffice. 2012. "OpenOffice.org." January 19. http://www.openoffice.org/distribution/p2p/

Ozer, Jan. 2011. "Streaming Vs. Progressive Download Vs. Adaptive Streaming." *Online Video.* May 13. http://www.onlinevideo.net/2011/05/streaming-vs-progressive-download-vs-adaptive-streaming/

P2PNet. 2009. "Vuze vs Cox the Copyright Cop." *P2PNet.* February http://www.p2pnet.net/story/18402

Parsons, Christopher. 2009. *Summary of January 13, 2009 CRTC Filings by Major ISPs in Response to Interrogatory PN 2008-19 with February 9, 2009 Updates.* http://www.christopher-parsons.com/PublicUpload/Summary_of_January_13_2009_ISP_filings_with_February_9_2009_Updates_version_1.0(for_web).pdf

Peterson, Larry L. & Bruce S. Davie. 2012. *Computer Networks: A Systems Approach.* 5th ed. Burlington, MA: Morgan Kaufmann.

Pierce, Richard J., Jr. . 2010. *Administrative Law Treatise.* 5th ed. Vol. I. New York, NY: Aspen Publishers.

Public Interest Advocacy Centre (PIAC), Canadian Association of Internet Providers (CAIP), Acanac Inc., Accelerated Connections Inc., Cybersurf Corp., eagle.ca, et al. 2009. Part VII Application to Review and Vary Telecom Decision CRTC 2008-108, Canadian Association of Internet Providers - Application Requesting Certain Orders Directing Bell Canada to Cease and Desist from Throttling its Wholesale ADSL Access Services to Canadian Radio-Television and Telecommunications Commission. File No. 8662-V42-200907826. May 21. http://crtc.gc.ca/public/partvii/2009/8662/p8_200907727/1140124.zip

Public Knowledge & Media Access Project. 2010. Notice of Oral Ex Parte Presentation to Federal Communications Commission. GN Dkt. No. 09-191. December 13. http://apps.fcc.gov/ecfs/document/view?id=7020923323

Quin, Julia Ya. 2005. "Defining Nondiscrimination Under the Law of the World Trade Organization." *Boston University International Law Journal*, 23: 215-297.

Quinn, Bob. 2010. "Who Keeps Pulling The Net Neutrality Football?" *AT&T Public Policy Blog.* September 9. http://attpublicpolicy.com/government-policy/who-keeps-pulling-the-net-neutrality-football/

Ranaweera, Chatura & Jaideep Prabhu. 2003. "The Influence of Satisfaction, Trust and Switching Barriers on Customer Retention in a Continuous Purchasing Setting." *International Journal of Service Industry Management*, 14(4): 374-395.

RCN Corporation. 2010. Ex Parte Letter to Federal Communications Commission. GN Dkt. No. 09-191. May 7. http://apps.fcc.gov/ecfs/document/view?id=7020450131

Read, Daniel. 2004. "Intertemporal Choice." In *Blackwell Handbook of Judgment and Decision Making,* eds. Nigel Harvey & Derek J. Koehler, 424-443. Oxford, UK: Blackwell Publishing.

Reardon, Marguerite 2007. "Comcast Denies Monkeying with BitTorrent Traffic." *CNet News.* August 21. http://www.news.com/8301-10784_3-9763901-7.html.

Reed, David. 2009. *Testimony of Doctor David Reed.* Attachment B Part 1 to the Campaign for Democratic Media's Comments In the Matter of Review of the Internet Traffic Management Practices of Internet Service Providers submitted February 23, 2009. File No. 8646-C12-200815400. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1029987.zip

Reed, David P. 2008. *Testimony before the U.S. House of Representatives, Subcommittee on Telecommunications and the Internet of the Committee on Energy and Commerce at its Hearing on: What Your Broadband Provider Knows About Your Web Use: Deep Packet Inspection and Communications Law and Policies.* http://www.gpo.gov/fdsys/pkg/CHRG-110hhrg58071/html/CHRG-110hhrg58071.htm

Reed, David P. 2010. "A Respone to Barbara van Schewick: Code Needs (Only a Little) Help From the Law." *dpr.* December 15. http://www.reed.com/blog-dpr/?p=85

Riley, M. Chris & Ben Scott. 2009. *Deep Packet Inspection: The End of the Internet as We Know It?* Free Press. White Paper. http://www.freepress.net/files/Deep_Packet_Inspection_The_End_of_the_Internet_As_We_Know_It.pdf

Riley, M. Chris & Robb Topolski. 2009. *The Hidden Harms of Application Bias.* Free Press and New America Foundation. Free Press/New America Foundation Policy Brief. http://newamerica.net/sites/newamerica.net/files/policydocs/The_Hidden_Harms_of_Application_Bias.pdf

Ripken, Susanna Kim. 2006. "The Dangers and Drawbacks of the Disclosure Antidote: Toward a More Substantive Approach to Securities Regulation." *Baylor Law Review*, 58(1): 139-204.

Rosenberg, Nathan. 1996. "Uncertainty and Technological Change." In *The Mosaic of Economic Growth,* eds. Ralph Landau, Timothy Taylor & Gavin Wright, 334-356. Stanford, CA: Stanford University Press.

Samuelson, William & Richard Zeckhauser. 1988. "Status Quo Bias in Decision Making." *Journal of Risk and Uncertainty*, 1(1): 7-59.

Sandvine. 2004. *Meeting the Challenge of Today's Evasive P2P Traffic: Service Provider Strategies for Managing P2P Filesharing.* Sandvine Inc.

Sandvine. 2005. *Turning Gaming into Revenue.* Sandvine, Inc. MSO Solution Overview. www.sandvine.com/general/getfile.asp?FILEID=86

Sandvine. 2013. *Success Story: Service Innovation with Third Party Partnerships.* https://www.sandvine.com/downloads/general/success-stories/success-story-vox-telecom-service-innovation-with-third-party-partnerships.pdf

Sandvine. 2014. *Global Internet Phenomena Report 1H 2014.* Sandvine. https://www.sandvine.com/downloads/general/global-internet-phenomena/2014/1h-2014-global-internet-phenomena-report.pdf

Schatz, Amy. 2008. "Cox About to Feel Wrath of Net Neutrality Activists." *WSJ Blogs: Washington Wire.* http://blogs.wsj.com/washwire/2008/05/15/cox-about-to-feel-wrath-of-net-neutrality-activists/

Schmidt, Sarah. 2012. "Complaints About Online Traffic Delays Accelerating, Says CRTC." *Canada.com.* January 12. http://www.canada.com/life/Complaints+about+online+traffic+delays+accelerating+says+CRTC/5986923/story.html

Schoen, Seth. 2007. "Comcast and BitTorrent." *EFF Deeplinks Blog.* September 13. http://www.eff.org/deeplinks/2007/09/comcast-and-bittorrent

Schuett, Florian. 2010. "Network Neutrality: A Survey of the Economic Literature." *Review of Network Economics*, 9(2).

Shalunov, Stanislav & Benjamin Teitelbaum. 2001. *QBone Scavenger Service (QBSS) Definition.* Internet2 Technical Report, Proposed Service Definition. http://qos.internet2.edu/wg/wg-documents/qbss-definition.txt

Shapiro, Carl & Hal R. Varian. 1998. *Information Rules. A Strategic Guide to the Network Economy.* Boston, MA: Harvard Business School Press.

Shaw, Russell. 2006. "Cable Broadband ISP's QoS Enhancement Surcharge Draws Vonage's Ire." *ZDNet IP Telephony Blog.* March 7. http://www.zdnet.com/blog/ip-telephony/cable-broadband-isps-qos-enhancement-surcharge-draws-vonages-ire/952

Shelanski, Howard. 2007. "Network Neutrality: Regulation with More Questions Than Answers." *Journal on Telecommunications and High Technology Law*, 6(1): 23-40.

Shields, Todd. 2010a. "AT&T Gains FCC's Ear as Regulators Near Decision on Net Neutrality Rules." *Bloomberg.* November 30. http://www.bloomberg.com/news/2010-11-30/at-t-gains-fcc-s-ear-as-regulators-near-decision-on-net-neutrality-rules.html

Shields, Todd. 2010b. "Critics Decry 'Secret Deal' as AT&T, Google Huddle With FCC." *Bloomberg.* August 3. http://www.bloomberg.com/news/2010-08-02/google-verizon-at-t-seek-internet-peace-in-closed-door-talks-with-fcc.html

Sidak, J. Gregory. 2006. "A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet." *Journal of Competition Law and Economics*, 2(3): 349-474.

Sidak, J. Gregory & David J. Teece. 2010. "Innovation Spillovers and the "Dirt Road" Fallacy: The Intellectual Bankruptcy of Banning Optional Transactions for Enhanced Delivery Over the Internet." *Journal of Competition Law & Economics*, 6(3): 521-594.

Silbey, Mari. 2012. "Stanford Team Wins Cable Show App Challenge." *SmartPlanet.* May 24. http://www.smartplanet.com/blog/thinking-tech/stanford-team-wins-cable-show-app-challenge/11743

Singel, Ryan. 2009. "Developer: Apple Denied Health Care App for Political Reasons." *Wired.* September 28. http://www.wired.com/epicenter/2009/09/apple-denied-health-care-app-for-political-reasons-developer-says/

Singer, Hal J. & J. Gregory Sidak. 2007. "Vertical Foreclosure in Video Programming Markets: Implications for Cable Operators." *Review of Network Economics*, 6(3): 372-395.

Skype. 2012. "Get Skype for Windows." February 21. http://www.skype.com/intl/en-us/get-skype/on-your-computer/windows/

Skype Communications S.A.R.L. 2010. Comments to Federal Communications Commission. GN Docket No. 09-191. January 14. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020377906

Solum, Lawrence B. & Minn Chung. 2004. "The Layers Principle: Internet Architecture and the Law." *Notre Dame Law Review*, 79(3): 815-947.

Soule, Helen, Mary Alice Baish, Michael McLeod-Bell, Lynne E. Bradley, Prudence Adler, Charles Benton, et al. 2009. *"Public Interest Letter in Support of FCC Public Proceedings to Ensure an Open and Nondiscriminatory Internet."* Letter to Julius Genachowski. October 21. http://www.publicknowledge.org/node/2711

Spangler, Todd. 2012. "Cable Show 2012: Stanford Team Wins 'App Challenge' With Bandwidth-Priority System." *Multichannel News.* May 23. http://www.multichannel.com/article/485029-Cable_Show_2012_Stanford_Team_Wins_App_Challenge_With_Bandwidth_Priority_System.php

Speta, James B. 2011. "Supervising Managed Services." *Duke Law Journal*, 60(8): 1715-1759.

Srinivasan, Venky & Vivek Gupta. 2010. Ex Parte Letter to Federal Communications Commission. GN Docket No. 09-191. December 10. http://apps.fcc.gov/ecfs/document/view?id=7020923207

St. Arnaud, Bill. 2009. *Report of Bill St. Arnaud.* Attachment C Part 1 to the Campaign for Democratic Media's Comments In the Matter of Review of the Internet Traffic Management Practices of Internet Service Providers submitted February 23, 2009. File No. 8646-C12-200815400. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1029987.zip

St. Johns, M. & Geoff Huston. 2003. "Considerations on the use of a Service Identifier in Packet Headers." Request for Comments 3639. IETF.

Stross, Randall. 2005. "Why Time Warner Has Fallen in Love With AOL, Again." *New York Times.* September 25, 3/3.

Sunstein, Cass R. 1995a. "Incompletely Theorized Agreements." *Harvard Law Review*, 108(7): 1733-1772.

Sunstein, Cass R. . 1995b. "Problems with Rules." *California Law Review*, 83(4): 953-1026.

Svensson, Peter 2007. Comcast Blocks Some Internet Traffic. *Associated Press.* October 19. http://www.sfgate.com/cgi-bin/article.cgi?f=/n/a/2007/10/19/financial/f061526D54.DTL

Teitelbaum, Ben & Stanislav Shalunov. 2002. "Why Premium IP Service Has Not Deployed (and Probably Never Will)." *Internet2 QoS Working Group Informational Document.* May 3, 2002 (HTMLized with updated references January 9, 2006). http://qos.internet2.edu/wg/documents-informational/20020503-premium-problems-non-architectural.html

Teitelbaum, Benjamin & Stanislav Shalunov. 2003. "What QoS Research Hasn't Understood About Risk." *ACM SIGCOMM Workshop on Revisiting IP QoS: What have we learned, why do we care?*: 148-150.

TELUS Communications Company. 2009. Comments to Canadian Radio-Television and Telecommunications Commission. File No. 8646-C12-200815400. February 23. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1029656.pdf

The Council of Scientific Society Presidents. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 13. http://apps.fcc.gov/ecfs/document/view?id=7020923181

The Hill Administrator. 2010. "Despite Seat at Table, AT&T Outed Delicate Net-Neutrality Talks." *The Hill.* November 25. http://thehill.com/blogs/hillicon-valley/technology/130737-despite-seat-at-the-table-atat-outed-delicate-net-neutrality-talks-sparking-gop-backlash

Thomas, Daniel. 2012. "Texting Revenues Hit by Web Services." *FT.* February 21. http://www.ft.com/intl/cms/s/0/5599fa26-5cb5-11e1-8f1f-00144feabdc0.html

U.S. Department of Justice. 2011. Competitive Impact Statement to United States District Court for the District of Columbia. United States of America, State of California, State of Florida, State of Missouri, State of Texas, and State of Washington v. Comcast Corp., General Electric Co. and NBC Universal, Inc. January 18. http://www.justice.gov/atr/cases/f266100/266158.pdf

Official Airline Guides v. Federal Trade Commission, 630 F.2d. 920. United States Court of Appeals, Second Circuit. 1980.

United States Government Accountability Office. 2006. *Report to Congressional Committees; Telecommunications; Broadband Deployment is Extensive throughout the United States, but It Is Difficult to Assess the Extent of Deployment Gaps in Rural Areas.* GAO-06-426. http://www.gao.gov/new.items/d06426.pdf

Van den Bossche, Peter. 2008. *The Law and Policy of the World Trade Organization: Text, Cases and Materials.* 2nd ed.: Cambridge University Press.

van Schewick, Barbara. 2004. "Architecture and Innovation: The Role of the End-to-End Arguments in the Original Internet." PhD Dissertation. Technical University Berlin.

van Schewick, Barbara. 2007. "Towards an Economic Framework for Network Neutrality Regulation." *Journal on Telecommunications and High Technology Law*, 5(2): 329-391.

van Schewick, Barbara. 2008a. *Official Testimony at the Federal Communications Commission Second En Banc Hearing on Broadband Management Practices.* Federal Communications Commission. www.fcc.gov/broadband_network_management/041708/vanschewick-written.pdf

van Schewick, Barbara. 2008b. *Oral Testimony at the Federal Communications Commission Second En Banc Hearing on Broadband Management Practices.* Federal Communications Commission. www.fcc.gov/broadband_network_management/041708/vanschewick-oral.pdf

van Schewick, Barbara. 2009. "Point/Counterpoint. Network Neutrality Nuances." *Communications of the ACM*, 52(2): 31-34, 36-37.

van Schewick, Barbara. 2010a. *Background Paper for the Federal Communications Commission's Workshop on Approaches to Preserving the Open Internet in Seattle, WA on April 28, 2010.*

van Schewick, Barbara. 2010b. "The FCC's Open Internet Proposal – Lessons from Silicon Valley." *Internet Architecture and Innovation Blog.* December 13. http://netarchitecture.org/2010/12/the-fccs-open-internet-proposal-lessons-from-silicon-valley/

van Schewick, Barbara. 2010c. "The FCC's Open Internet Rules – Stronger than You Think." *Internet Architecture and Innovation Blog.* December 27. http://netarchitecture.org/2010/12/the-fcc's-open-internet-rules-–-stronger-than-you-think/

van Schewick, Barbara. 2010d. *Internet Architecture and Innovation.* Cambridge, MA: MIT Press.

van Schewick, Barbara. 2010e. *Network Neutrality: What a Non-Discrimination Rule Should Look Like.* Attachment to Barbara van Schewick's Notice of Ex Parte Conversation In the Matter of Preserving the Open Internet submitted August 2, 2010. GN Dkt. No 09-191. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020652518

van Schewick, Barbara. 2010f. *Network Neutrality: What a Non-Discrimination Rule Should Look Like (Version 1.0). Center for Internet and Society White Paper.* Attachment to Barbara van Schewick's Ex Parte Letter In the Matter of Preserving the Open Internet submitted December 14, 2010. GN Dkt. No 09-191. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020923837

van Schewick, Barbara. 2010g. *Network Neutrality: What a Non-Discrimination Rule Should Look Like. Paper presented at the 38th Research Conference on Communication, Information and Internet Policy (TPRC 2010).* Stanford Public Law Working Paper No. 1684677. Stanford Law and Economics Olin Working Paper No. 402. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1684677

van Schewick, Barbara. 2010h. *Opening Statement at the Federal Communications Commission's Workshop on Approaches to Preserving the Open Internet.* Federal Communications Commission. http://www.law.stanford.edu/display/images/dynamic/publications_pdf/schewick-statement-20100428.pdf

van Schewick, Barbara. 2010i. *Opening Statement at the Federal Communications Commission's Workshop on Innovation, Investment and the Open Internet.* Federal Communications Commission. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020382126

van Schewick, Barbara. 2010j. *Outline of Proposal for Open Internet Rules.* Attachment to Barbara van Schewick's Notice of Ex Parte Conversations In the Matter of Preserving the Open Internet submitted on May 27, 2010. GN Dkt. No 09-191. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020493599

van Schewick, Barbara. 2010k. Notice of Ex Parte Meeting to Federal Communications Commission. GN Dkt. No 09-191. December 9. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020922658

van Schewick, Barbara. 2010l. "Start-Up Video Company Asks FCC to Improve Open Internet Proposal." *Internet Architecture and Innovation Blog.* December 12. http://netarchitecture.org/2010/12/start-up-video-company-files-concerns-about-fcc-open-internet-proposal/

van Schewick, Barbara. 2011a. "Is Verizon Wireless Illegally Blocking Google Wallet? It's Time for the FCC to Investigate." *Internet Architecture and Innovation Blog.* December 19. https://netarchitecture.org/2011/12/is-verizon-wireless-illegally-blocking-google-wallet-its-time-for-the-fcc-to-investigate/

van Schewick, Barbara. 2011b. "Public Interest Requires Public Input: Verizon/Android Tethering." *Internet Architecture and Innovation Blog.* June 30. https://netarchitecture.org/2011/06/public-interest-requires-public-input-verizonandroid-tethering/

van Schewick, Barbara. 2012a. Comments to European Commission's Public Consultation on Specific Aspects of Transparency, Traffic Management and Switching in an Open Internet. October 15

van Schewick, Barbara. 2012b. *Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like.* Center for Internet and Society White Paper. http://cyberlaw.stanford.edu/publications/network-neutrality-and-quality-service-what-non-discrimination-rule-should-look

van Schewick, Barbara. 2012c. "Network Non-Discrimination and Quality of Service." *Internet Architecture and Innovation Blog.* June 13. http://netarchitecture.org/2012/06/network-non-discrimination-and-quality-of-service-2/

van Schewick, Barbara. 2014a. "The Case for Rebooting the Network-Neutrality Debate." *The Atlantic.* May 6. http://www.theatlantic.com/technology/archive/2014/05/the-case-for-rebooting-the-network-neutrality-debate/361809/

van Schewick, Barbara. 2014b. "The FCC Changed Course on Network Neutrality. Here is Why You Should Care." *Stanford Center for Internet and Society Blog.* April 25. http://cyberlaw.stanford.edu/blog/2014/04/fcc-changed-course-network-neutrality-here-why-you-should-care

Varian, Hal R. 1999. *Intermediate Microeconomics. A Modern Approach.* 5th ed. New York: W. W. Norton & Company, Inc.

Vaxination Informatique. 2009. Application to Review and Vary Telecom Decision CRTC 2008-108, Canadian Association of Internet Providers - Application Requesting Certain Orders Directing Bell Canada to Cease and Desist from Throttling its Wholesale ADSL Access Services to Canadian Radio-Television and Telecommunications Commission. File No. 8662-P8-200907727. May 20. http://crtc.gc.ca/public/partvii/2009/8662/v42_200907826/1148017.zip

Verizon. 2006. *Verizon Private IP Presentation.* http://www.verizonbusiness.com/resources/presentations/pr_private-ip-presentation_en_xg.pdf

Verizon. 2010. *Verizon Private IP Service Differentiators.* Verizon. http://www22.verizon.com/wholesale/attachments/solutions/W5036.pdf

Verizon & Verizon Wireless. 2007. Comments to Federal Communications Commission. WC Dkt. No. 07-52. June 15. http://apps.fcc.gov/ecfs/document/view?id=6519529411

Verizon & Verizon Wireless. 2010. Comments to Federal Communications Commission. GN Docket No. 09-191. January 14. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020378523

von Bornstaedt, Falk. 2012. *Deutsche Telekom Group - Quality-of-Service (QoS) for IP-Transit / Peering.* White Paper Version 1.02. https://bscw.ict-etics.eu/pub/bscw.cgi/d44579/DT%20QoS%20White%20Paper.pdf

Waxman, Henry A. 2010. "Waxman Statement on Net Neutrality Proposal." September 29. http://waxman.house.gov/chairman-waxmans-statement-net-neutrality-proposal

Weiser, Philip J. 2003. "Toward a Next Generation Regulatory Strategy." *Loyola University of Chicago Law Journal*, 35(1): 41-85.

Weiser, Philip J. 2008. "The Next Frontier for Network Neutrality." *Administrative Law Review*, 60(2): 273-322.

Werner, Tony. 2012. "The Facts about Xfinity TV and Xbox 360: Comcast is Not Prioritizing." *Comcast Voices.* May 15. http://blog.comcast.com/2012/05/the-facts-about-xfinity-tv-and-xbox-360-comcast-is-not-prioritizing.html

Wernick, Christian. 2007. *Strategic Investment Decisions in Regulated Markets. The Relationship Between Infrastructure Investments and Regulation in European Broadband.* Deutscher Universitäts-Verlag.

Wilson, Fred. 2010. "Regulation Strangulation." *AVC.* August 12. http://www.avc.com/a_vc/2010/08/regulation-strangulation.html

Wise, Andrew Stewart & Kiran Duwadi. 2005. "Competition Between Cable Television and Direct Broadcast Satellite: The Importance of Switching Costs and Regional Sports Networks." *Journal of Competition Law and Economics*, I(4): 679-705.

Wu, Tim. 2003. "Network Neutrality and Broadband Discrimination." *Journal on Telecommunications & High Technology Law*, 2: 141-175.

Wu, Tim. 2004. "The Broadband Debate: A User's Guide." *Journal on Telecommunications & High Technology Law*, 3(1): 69-95.

Wu, Tim. 2006a. "The AT&T Network Neutrality Agreement." December 29. http://www.timwu.org/log/archives/81

Wu, Tim. 2006b. *Testimony before the House Committee on the Judiciary: Telecom & Antitrust Task Force on Network Neutrality: Competition, Innovation, and Nondiscriminatory Access.* 109th Congress, 2nd Session.

Wu, Tim. 2006c. "Why Have a Telecommunications Law? Anti-Discrimination Norms in Communications." *Journal on Telecommunications and High Technology Law*, 5: 15-46.

Wu, Tim. 2006d. "Why You Should Care About Network Neutrality." *Slate.* May 1. http://www.slate.com/id/2140850

Wu, Tim & Lawrence Lessig. 2003. Ex Parte Submission to Federal Communications Commission. CS DKT. No. 02-52. August 22. http://fjallfoss.fcc.gov/ecfs/document/view?id=6514683885

Yiakoumis, Yiannis. 2013. *"Questions on your work on user-controlled QoS."* Email to Barbara van Schewick. March 20.

Yiakoumis, Yiannis, Te-Yuan Huang, Kok-Kiong Yap, Sachin Katti, Nick McKeown & Ramesh Johari. 2012. "Putting Home Users in Charge of their Network." *2012 ACM Ubicomp Workshop on Systems and Infrastructure for the Digital Home, HomeSys'12*: 1114-1119.

Yiakoumis, Yiannis, Kok-Kiong Yap, Sachin Katti, Guru Parulkar & Nick McKeown. 2011. "Slicing Home Networks." Paper presented at 2011 ACM Sigcomm Workshop on Home Networks, HomeNets '11. Toronto, Canada.

Yoo, Christopher. 2005. "Beyond Network Neutrality." *Harvard Journal of Law and Technology*, 19(1): 1-77.

Yoo, Christopher. 2006. "Network Neutrality and the Economics of Congestion." *Georgetown Law Journal*, 94(June): 1847-1908.

Yoo, Christopher. 2007. "What Can Antitrust Contribute to the Network Neutrality Debate?" *International Journal of Communication*, 1: 493-530.

Yoo, Christopher. 2008. "Network Neutrality, Consumers, and Innovation." *University of Chicago Legal Forum*, 2008: 179-262.

Yoo, Christopher. 2009. "Network Neutrality after Comcast: Toward a Case-by-Case Approach to Reasonable Network Management." In *New Directions in Communications Policy,* ed. Randolph J. May, 55-84. Durham, NC: Carolina Academic Press.

Zachem, Kathryn A. 2009. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. January 30. http://fjallfoss.fcc.gov/ecfs2/document/view.action?id=6520212400

Ziegler, Chris. 2012. "AT&T Only Allowing FaceTime Over Cellular on Mobile Share Plans, No Extra Charge." *The Verge.* August 17. http://www.theverge.com/2012/8/17/3250228/att-facetime-over-cellular-ios-6-mobile-share