ARTICLES

NETWORK NEUTRALITY AND QUALITY OF SERVICE: WHAT A NONDISCRIMINATION RULE SHOULD LOOK LIKE

Barbara van Schewick*

1

Over the past ten years, the debate over "network neutrality" has remained one of the central debates in Internet policy. Governments all over the world have been investigating whether legislative or regulatory action is needed to limit the ability of providers of Internet access service to interfere with the applications, content, and services on their networks.

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. Nondiscrimination rules apply to any form of differential treatment that falls short of blocking. Policymakers who consider adopting network neutrality rules need to decide which, if any, forms of differential treatment should be banned.

This Article makes five contributions: First, it proposes a substantive framework that policymakers can use to evaluate alternative proposals for network neutrality rules and assess specific forms of discriminatory conduct. Second, the Article evaluates eight existing proposals for nondiscrimination rules and the Open Internet Order's nondiscrimination rule against this framework and proposes a nondiscrimination rule—ban application-specific discrimination, allow application-agnostic discrimination—that policymakers should adopt around the world—a rule that the FCC's Open Internet Order adopted in part. Third, the Article 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 Article offers the first in-depth analysis of the relationship between network neutrality and new network-level services called Quality of Service. Finally, the Article provides the first detailed analysis of the Open Internet Order's nondiscrimination rule for fixed broadband Internet access.

The questions this Article addresses are global, and not limited to American law and policy: In Europe, the European Commission, the European Parliament, and the member states are currently considering which approach to network neutrality they should take. In Brazil, policymakers are discussing the best way to implement the recently adopted network neutrality law. In all of these debates, nondiscrimination 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.

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 should Internet users decide? 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. Initially, the network was application-blind: it could not distinguish between the applications, content, and services that were running over the network.[2] As a

---

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

    2. Throughout this Article, the term "original architecture of the Internet" refers to the network architecture that was specified in the Defense Advanced Research Projects Agency (DARPA) Internet Program Protocol Specifications for the Internet Protocol, INFO. SCIS. INST., UNIV. OF S. CAL., RFC 791, INTERNET PROTOCOL: DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION (Jon Postel ed., 1981), and Transmission Control Protocol, INFO. SCIS. INST., UNIV. OF S. CAL., RFC 793, TRANSMISSION CONTROL PROTOCOL: DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION (Jon Postel ed., 1981). David Clark described this architecture in an important article on the design philosophy of the DARPA Internet protocols. *See* David D. Clark, *The Design Philosophy of the DARPA Internet Protocols*, COMPUTER COMM. REV., Aug. 1988, at 106. 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. BARBARA VAN SCHEWICK, INTERNET ARCHITECTURE AND INNOVATION 57-81, 377-79 (2010) [hereinafter VAN SCHEWICK, ARCHITECTURE AND INNOVATION]; Barbara van Schewick, Architecture and Innovation: The Role of the End-to-End Arguments in the Original Internet 87-109, 123-29 (Sept. 15, 2004) (unpublished Ph.D. dissertation, Technical University Berlin) (on file with author) [hereinafter van Schewick, Dissertation]. Both versions shaped the original architecture of the Internet. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra*, at 90-103, 110-12, 379-81; van Schewick, Dissertation, *supra*, at 116-29. However, only the broad version, together with the layering principle,

result, Internet service 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 Internet service providers. Over the past two decades, technology has become available that enables Internet service providers to identify the applications and content on their networks and control their execution.[3]

Proponents of network neutrality argue that Internet service providers have incentives to use this new technology in socially harmful ways.[4] 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. 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 as the debate over network neutrality. The Open Internet proceeding, started by the Federal Communications Commission (FCC) in the fall of 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*. As of October 2014, more than 3.9 million comments had been filed in the FCC's current network neutrality rulemaking.[5] All over the world, from the United States to Europe to Latin America, policymakers continue to investigate whether they should adopt network neutrality rules and, if so, what the rules should be.

---

is responsible for the application-blindness of the network. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra*, at 72-75, 217-18; van Schewick, Dissertation, *supra*, at 101-03; *see also* David P. Reed, *A Response to Barbara van Schewick: Code Needs (Only a Little) Help from the Law*, DPR (Dec. 15, 2010), http://www.reed.com/blog-dpr/?p=85. On the layering principle and its relationship to the architecture of the Internet, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra*, at 46-57, 88-90.

3. On deep packet inspection (DPI) in general, see Nate Anderson, *Deep Packet Inspection Meets 'Net Neutrality, CALEA*, ARS TECHNICA (July 25, 2007, 9:10 PM PDT), http://arstechnica.com/gadgets/2007/07/deep-packet-inspection-meets-net-neutrality. For a specific example, see *Network-Based Application Recognition and Distributed Network-Based Application Recognition*, CISCO SYS., http://www.cisco.com/c/en/us/td/docs/ios/12 _2s/feature/guide/fsnbarad.pdf (last visited Jan. 7, 2015). On the state of DPI deployment, see Comments of Free Press at 141-51, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Free Press Open Internet Comments], *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id =7020378751.

4. This paragraph is adapted from Barbara van Schewick & David Farber, *Point/Counterpoint: Network Neutrality Nuances*, COMM. ACM, Feb. 2009, at 31, 32.

5. Gigi B. Sohn, *FCC Releases Open Internet Reply Comments to the Public*, OFFICIAL FCC BLOG (Oct. 22, 2014, 3:07 PM), http://www.fcc.gov/blog/fcc-releases-open -internet-reply-comments-public.

In Europe, the European Commission, the European Parliament, and the member states are currently considering which approach to network neutrality they should take.[6] In Brazil, policymakers are discussing the best way to implement the recently adopted network neutrality law.[7] In the United States, a 2014 decision by the Court of Appeals for the D.C. Circuit reignited the debate. In December 2010, the FCC adopted the Open Internet Order,[8] which enacted binding network neutrality rules for the first time.[9] The rules went into effect in November 2011.[10] In January 2014, *Verizon v. FCC* struck down the core provisions of that Order—the rules against blocking and discrimination.[11] 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."[12] Both of these points had been heavily contested by Verizon. The court found, however, that the no-blocking and nondiscrimination rules violated 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.[13] 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 require 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, U.S. policymakers must decide (again) which, if any, network neutrality rules the United States should adopt. They essentially have three options. First, the FCC can preserve the Open Internet Rules by reclassifying Internet service as a telecommunications service

---

6. *See, e.g.*, James Kanter & Mark Scott, *European Panel Adopts 'Net Neutrality' and Mobile Roaming Rules*, N.Y. TIMES (Mar. 18, 2014), http://www.nytimes.com/2014 /03/19/business/international/eu-panel-adopts-net-neutrality-and-mobile-roaming-rules.html.

7. *See, e.g.*, *Net Neutrality Wins in Brazil's 'Internet Constitution,'* AL JAZEERA AM. (Mar. 26, 2014, 9:34 AM ET), http://america.aljazeera.com/articles/2014/3/26/brazil-internet -constitution.html.

8. Preserving the Open Internet (*Open Internet Order*), 25 FCC Rcd. 17,905 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

9. *Id.* at 17,908-15.

10. Preserving the Open Internet, 76 Fed. Reg. 59,192 (Sept. 23, 2011) (codified at 47 C.F.R. pts. 0, 8).

11. *Verizon*, 740 F.3d at 659.

12. *Id.* at 628.

13. *Id.* at 655-56.

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 D.C. Circuit'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 nondiscrimination rule—and, if so, what it should be—is a key point of contention in all of these debates. This Article analyzes the available options and proposes a nondiscrimination rule—ban application-specific discrimination, allow application-agnostic discrimination—that policymakers should adopt around the world—a rule that the FCC's Open Internet Order adopted in part.

Nondiscrimination 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 applications, 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.[14]

The decision for a specific nondiscrimination 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.[15]

---

14. An Internet service provider's pricing practices with respect to application and content providers who are not its Internet service customers are not governed by the nondiscrimination rule. Whether an Internet service provider can charge application providers who are not its subscribers an "access fee" for access or for prioritized or otherwise enhanced access to its subscribers is governed by a network neutrality regime's rules about access fees. For a definition of the term "access fees," see note 62 below. A discussion of access fees is outside the scope of this Article. For a short overview of the debate, see Box 2 and notes 29-30 below.

15. For example, Internet telephony is very sensitive to delay above a certain level, but can tolerate occasional packet loss. Users usually do not notice a one-way, mouth-to-ear delay of less than 150 milliseconds (ms). A delay of more than 400 ms makes voice calls frustrating or unintelligible. *See* INT'L TELECOMM. UNION, RECOMMENDATION G.114, ONE-WAY TRANSMISSION TIME 3 (2003); JAMES F. KUROSE & KEITH W. ROSS, COMPUTER NETWORKING: A TOP-DOWN APPROACH 601 (5th ed. 2010). Depending on the encoding and loss-concealment mechanisms used, Internet telephony applications can tolerate between 1% and 20% of packet loss. *See* KUROSE & ROSS, *supra*, at 617. By contrast, e-mail is very sensitive to packet loss, but can tolerate some delay. *Id.* at 92, 95 fig.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, for example, *id.* at 92-95; and LARRY L. PETERSON & BRUCE S. DAVIE, COMPUTER NETWORKS: A SYSTEMS APPROACH 530-37 (5th ed. 2012).

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),[16] a network that provides Quality of Service offers different types of service to different data packets.[17] 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.[18] 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.[19]

Thus, policymakers who consider adopting nondiscrimination rules face a serious challenge: how to find a nondiscrimination rule that realizes the goals of network neutrality regulation without overly constraining the evolution and operation of the network and 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 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 nondiscrimination 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 Article addresses that challenge. It makes five contributions[20]:

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 noneconomic ways. A lot of uncertainty exists, however, on how to move from

---

16. Thus, the network operates like the default service offered by the U.S. Postal Service, which does not guarantee when a letter will arrive or whether it will arrive at all. In contrast to the Postal Service, which lets users choose services other than the default service, such as two-day shipping, the original Internet provides only best-effort service. PETERSON & DAVIE, *supra* note 15, at 206-07.

17. *See, e.g.*, KUROSE & ROSS, *supra* note 15, at 647-72 (discussing quality of service); PETERSON & DAVIE, *supra* note 15, at 530-57 (same).

18. For example, of the two Quality of Service architectures that were standardized by the Internet Engineering Task Force (IETF), the IntServ architecture provides specific Quality of Service guarantees to particular application sessions, while the DiffServ architecture provides different levels of performance to different classes of traffic without providing specific guarantees. *See* KUROSE & ROSS, *supra* note 15, at 669-72 (discussing IntServ); *id.* at 660-65 (discussing DiffServ).

19. For a more detailed discussion of Quality of Service and of the conditions under which it is useful, see Part II.B.2 below.

20. The following discussion draws in part on Barbara van Schewick, *Network Non-Discrimination and Quality of Service*, INTERNET ARCHITECTURE & INNOVATION (June 13, 2012, 7:00 AM), http://netarchitecture.org/2012/06/network-non-discrimination-and-quality-of-service-2.

a high-level commitment to network neutrality to a concrete set of rules. This Article proposes a framework that policymakers 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. These factors are user choice, application-agnosticism, innovation without permission, and low costs of application innovation. Network neutrality rules should make it easy to determine which behavior is and is not allowed to provide much-needed certainty for industry participants, should keep the costs of regulation low, and should avoid constraining the evolution of the network more than is necessary to reach these goals.

Second, over the years, the network neutrality debate has evolved into a series of subdebates that are difficult to follow. As a result, even interested parties often lack a complete, accurate picture of the broader debate. This Article provides a comprehensive overview of existing proposals for nondiscrimination rules from a variety of sources, including academics, industry participants, and policymakers in the United States and abroad, and helps policymakers choose among the available options. It evaluates these proposals and the Open Internet Order's nondiscrimination rule against the framework developed in Part I and proposes a nondiscrimination rule that policymakers 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 have 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).

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.

This Article also explains why the following proposals for nondiscrimination rules do not adequately protect the values that network neutrality rules are

designed to protect: rules that allow all discrimination; rules that ban 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; and 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 United States or abroad.

Third, this Article exposes the deep disconnect between those, including the FCC in the Open Internet Order, who base calls for network neutrality regulation on a broad theoretical framework that considers a wide range of economic and noneconomic harms and those who evaluate calls for network neutrality regulation based on an antitrust framework. As this Article 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[21]—which is correct if you view the debate through an antitrust framework. In the context of the broad theoretical framework that the FCC explicitly adopted 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, in many cases, 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, this Article 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 nondiscrimination rule affects network providers' ability to offer Quality of Service.

The network neutrality debate is often framed as a debate for or against Quality of Service.[22] As this Article shows, the reality is much more subtle.

---

21. *See, e.g.*, Thomas W. Hazlett & Joshua D. Wright, *The Law and Economics of Network Neutrality*, 45 IND. L. REV. 767, 771-72, 796-803 (2012).

22. This perception is particularly common in Europe. *See, e.g.*, Martin Cave & Pietro Crocioni, *Does Europe Need Network Neutrality Rules?*, 1 INT'L J. COMM. 669, 671 (2007) (calling "no prioritisation" one of "the two main elements . . . of net neutrality"); Florian Schuett, *Network Neutrality: A Survey of the Economic Literature*, 9 REV. NETWORK ECON., no. 2, 2010, at 1, 1 ("Roughly speaking, network neutrality refers to the principle that all data packets on an information network are treated equally."); Martin Cave et al., Statement by

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.[23] Underlying these differences are disagreements over the social desirability of different forms of Quality of Service. This Article 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 nondiscrimination rule proposed by this Article and the Open Internet Order's nondiscrimination rule allow network providers to offer these (and only these) forms of Quality of Service.[24]

Finally, this Article provides a detailed analysis of the Open Internet Order's nondiscrimination rule for fixed broadband Internet access and of its implications for network providers' ability to manage their networks and offer Quality of Service. An accurate understanding of the rule is central to the current debate over the future of network neutrality in the United States. The rule bans discrimination that is "unreasonable," subject to reasonable network management.[25] 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

---

European Academics on the Inappropriateness of Imposing Increased Internet Regulation in the EU 1 (Jan. 8, 2009), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id =1329926 ("They [sic] key issue is whether internet service providers should be prevented from introducing differentiated quality of service levels on the Internet."); AT&T et al., *Ensuring Network Stability and Consumer Confidence in Competitive Markets*, LA QUADRATURE DU NET 2 (2008), http://www.laquadrature.net/files/net.confidence.coalition .pdf ("'Net neutrality' advocates typically suggest . . . an 'all bits must be treated the same' approach.").

   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, ARCHITECTURE AND INNOVATION, *supra* note 2, at 106-07.

   23. Quality of Service can be offered in different ways that 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.

   24. This Article's findings with respect to network neutrality and Quality of Service are summarized in more detail in the Conclusion.

   25. 47 C.F.R. § 8.7 (2014), *invalidated by* Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

reading of the text of the Order, the Article sets out the Open Internet Order's nondiscrimination 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.

According to the text of the Order, the FCC was to evaluate discriminatory conduct under the nondiscrimination rule and the reasonable network management exception based on how well the conduct preserves two of the factors that were at the core of the Internet's success: user choice and application-agnosticism.[26] This standard allows certain forms of user-controlled Quality of Service. The Order explicitly rejects attempts to base nondiscrimination 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 nondiscrimination rule proposed by this Article, the proposal heavily influenced the Open Internet Order's nondiscrimination 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 was to use to determine whether the conduct violates the FCC's nondiscrimination rule and the reasonable network management exception. Thus, this Article's discussion of application-specific and application-agnostic discrimination can illuminate the rationale underlying the FCC's rule and illustrate how these provisions apply to specific instances of discriminatory conduct.

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

This Article 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 subdebates. Each subdebate 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 below.)[27]

---

26. In addition, preserving the freedom to innovate without permission is an explicit purpose of the Open Internet Rules. 47 C.F.R. § 8.1. Thus, this factor can be used to interpret any provision of the Rules, including the nondiscrimination rule.

27. Throughout this Article, boxes are used to provide additional information that is relevant to the argument in the text without interrupting the flow of the main argument. They

The first question is as follows: Do we need a rule against blocking, that is, 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 Article assumes that the case for a rule against blocking has been made.[28]

This Article 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 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 Article 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. **Goal of the Rules**
    General-purpose infrastructure vs. antitrust

B. **Rules**
    1. Ban blocking?
    2. *Ban discrimination? (subject of this Article)*
    3. Impose limits on charging?
        i. for access to end users?
        ii. for enhanced access to end users?

C. **Exceptions**
    1. For what?
    2. How to define?

---

may, for example, explain a concept that some readers may not be familiar with, or explore an issue in more depth.

    28. *See, e.g.*, VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, chs. 6-9, at 215-375.

> **D. Scope**
>   1. Wireline vs. wireless
>   2. Internet access services vs. specialized services

This Article 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 have taken up restrictions on charging elsewhere.[29] There, I argue that network providers should only be allowed to charge their own Internet service customers for any differential treatment allowed by the nondiscrimination rule.[30]

> BOX 2
>
> CHARGING FOR QUALITY OF SERVICE
>
> If policymakers adopt a nondiscrimination 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, policymakers have a number of options, each supported by at least some proponents of network neutrality: (1) the

---

29. *See id.* at 278-80, 290-93; Barbara van Schewick, Assistant Professor of Law, Stanford Law Sch., Opening Statement at the Federal Communications Commission's Workshop on Approaches to Preserving the Open Internet (Apr. 28, 2010) [hereinafter van Schewick, Open Internet Opening Statement], *available at* https://www.law.stanford.edu /sites/default/files/publication/259136/doc/slspublic/schewick-statement-20100428.pdf; Barbara van Schewick, Outline of Presentation for the Federal Communications Commission's Workshop on Approaches to Preserving the Open Internet 10-12 (2010) (unpublished background paper) (on file with author) [hereinafter van Schewick, Background Paper]; Barbara van Schewick, *The Case for Rebooting the Network-Neutrality Debate*, ATLANTIC (May 6, 2014, 2:37 PM ET), http://www.theatlantic.com/technology/archive/2014/05/the-case-for -rebooting-the-network-neutrality-debate/361809 [hereinafter van Schewick, *Rebooting the Network-Neutrality Debate*]; Barbara van Schewick, *The FCC Changed Course on Network Neutrality. Here Is Why You Should Care.*, STAN. LAW SCH. CTR. FOR INTERNET & SOC'Y BLOG (Apr. 25, 2014, 7:16 AM), http://cyberlaw.stanford.edu/blog/2014/04/fcc-changed -course-network-neutrality-here-why-you-should-care [hereinafter van Schewick, *FCC Changed Course*].

30. 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, Open Internet Opening Statement, *supra* note 29; van Schewick, Background Paper, *supra* note 29, at 7-10.

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);[31] (2) it can charge only its Internet service customers;[32] (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 nondiscriminatory basis;[33] or (4) it can charge its Internet service customers and/or application and content providers.[34]

In sum, this Article 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.

---

31. *See, e.g.*, Internet Freedom Preservation Act, S. 215, 110th Cong. § 2 (2007); Internet Freedom and Nondiscrimination Act of 2006, H.R. 5417, 109th Cong. § 3 (2006); Network Neutrality Act of 2006, H.R. 5273, 109th Cong. § 4(a)(7) (2006).

32. *See, e.g.*, *Network Neutrality: Hearing Before the S. Comm. on Commerce, Sci. & Transp.*, 109th Cong. 57-58 (2006) [hereinafter *Hearing on Network Neutrality*] (prepared statement of Lawrence Lessig, Professor of Law, Stanford Law School). For criticism of this proposal, see Susan P. Crawford, *The Internet and the Project of Communications Law*, 55 UCLA L. REV. 359, 403-04 (2007); and Brett M. Frischmann, *An Economic Theory of Infrastructure and Commons Management*, 89 MINN. L. REV. 917, 1009-12 (2005). 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, S. 2360, 109th Cong. § 4(a)(3) (2006); Robin S. Lee & Tim Wu, *Subsidizing Creativity Through Network Design: Zero-Pricing and Net Neutrality*, J. ECON. PERSP., Summer 2009, at 61, 63-64; van Schewick, Open Internet Opening Statement, *supra* note 29, at 1; van Schewick, Background Paper, *supra* note 29, at 7-10; van Schewick, *Rebooting the Network-Neutrality Debate*, *supra* note 29; van Schewick, *FCC Changed Course*, *supra* note 29. For criticism of these proposals, see Robert Hahn & Scott Wallsten, *The Economics of Net Neutrality*, ECONOMISTS' VOICE, June 2006, at 1, 3-5; and C. Scott Hemphill, *Network Neutrality and the False Promise of Zero-Price Regulation*, 25 YALE J. ON REG. 135, 145-50 (2008).

33. *See, e.g.*, *The Future of the Internet: Hearing Before the S. Comm. on Commerce, Sci. & Transp.*, 110th Cong. 56 (2008) [hereinafter *Hearing on the Future of the Internet*] (prepared statement of Lawrence Lessig, Professor of Law, Stanford Law School). For criticism of this proposal, see J. Gregory Sidak, *A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet*, 2 J. COMPETITION L. & ECON. 349, 349-50 (2006).

34. *See, e.g.*, Comments of Verizon & Verizon Wireless at 68-74, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Verizon Open Internet Comments], *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=7020378523; Comments of Verizon & Verizon Wireless on the Notice of Inquiry at 35-38, Broadband Industry Practices, WC Docket No. 07-52 (June 15, 2007) [hereinafter Verizon Broadband Industry Practices Comments], *available at* http://apps.fcc.gov/ecfs/document/view?id=6519529411.

The Article proceeds in three parts. Part I sets out criteria that policymakers 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 nondiscrimination rules against these criteria and proposes a nondiscrimination rule—ban application-specific discrimination, allow application-agnostic discrimination—that policymakers should adopt. In the process, it explains how the different nondiscrimination rules affect network providers' ability to offer Quality of Service and which, if any, forms of Quality of Service a nondiscrimination rule should allow.

Part III sets out the Open Internet Order's nondiscrimination rule for fixed broadband Internet access. It evaluates the rule against the criteria used throughout this Article and discusses how the rule would affect network providers' ability to offer Quality of Service.

Opponents of network neutrality regulation have created the impression that network neutrality rules force policymakers 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 Article refutes that myth. It shows how policymakers can protect users and innovators while also giving network providers the tools they need to manage their networks and allowing the network to evolve.

I. A FRAMEWORK FOR EVALUATING NETWORK NEUTRALITY RULES

When evaluating alternative proposals for nondiscrimination or other network neutrality rules, legislators or regulators should consider a number of factors. Nondiscrimination 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.[35] 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 noneconomic harms.[36] The FCC's Open Internet Rules are based on this broader framework

---

35. *See* sources cited *infra* Part II.C.1.

36. *See, e.g.*, *Network Neutrality: Competition, Innovation, and Nondiscriminatory Access: Hearing Before the Task Force on Telecom & Antitrust of the H. Comm. on the Judiciary*, 109th Cong. 57-59 (2006) [hereinafter *Task Force Hearing*] (prepared statement of

as well.[37] Due to these differences, proponents of an antitrust framework and proponents of a broader framework will reach differing conclusions when evaluating proposals.[38] This Article assumes that the case for the broader theoretical framework has been made.

According to this broader theoretical framework, network neutrality regulation serves three major 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 time in various economic and noneconomic ways.[39] More specifically, network neutrality rules aim, first, to foster innovation in applications.[40] Fostering application innovation not only is

---

Timothy Wu, Professor of Law, Columbia Law School); Comments of the Center for Democracy & Technology at 8-9, 30, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Center for Democracy & Technology Comments], *available at* http://fjallfoss.fcc.gov/ecfs/document /view?id=7020378292; Free Press Open Internet Comments, *supra* note 3, at 76-77, 134-36; Comments of the Open Internet Coalition at 76, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Open Internet Coalition Comments], *available at* http://apps.fcc.gov/ecfs/document /view?id=7020377928; Comments of Public Interest Commenters at 23, 24-28, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Public Interest Comments], *available at* http://apps.fcc .gov/ecfs/document/view?id=7020378818; Crawford, *supra* note 32, at 380-92; Susan P. Crawford, *Transporting Communications*, 89 B.U. L. Rev. 871, 916-17, 919 (2009); Frischmann, *supra* note 32, at 1012-22; Brett M. Frischmann & Barbara van Schewick, *Network Neutrality and the Economics of an Information Superhighway: A Reply to Professor Yoo*, 47 Jurimetrics J. 383, 423-28 (2007); Lawrence Lessig, *The Internet Under Siege*, Foreign Pol'y (Nov. 1, 2001), https://web.archive.org/web/20120129230040/http://www .foreignpolicy.com/articles/2001/11/01/the_internet_under_siege (accessed via the Internet Archive index).

37. *See* 47 C.F.R. § 8.1 (2014); *Open Internet Order*, 25 FCC Rcd. 17,905, 17,908-15, 17,949-50 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

38. *See, e.g.*, Frischmann & van Schewick, *supra* note 36, at 426-28; Jerry Kang, *Race.Net Neutrality*, 6 J. on Telecomm. & High Tech. L. 1, 6-14 (2007); *see also infra* Part II.C.1.

39. *See, e.g.*, van Schewick, Architecture and Innovation, *supra* note 2, at 387-88; Frischmann & van Schewick, *supra* note 36, at 423-28; *see also* 25 FCC Rcd. at 17,908-15, 17,949-50.

40. *See, e.g.*, Lawrence Lessig, The Future of Ideas: The Fate of the Commons in a Connected World 246-47 (2001); Barbara van Schewick, *Towards an Economic Framework for Network Neutrality Regulation*, 5 J. on Telecomm. & High Tech. L. 329, 332 (2007); Tim Wu, *The Broadband Debate, A User's Guide*, 3 J. on Telecomm. & High Tech. L. 69, 72-74, 80-84 (2004) [hereinafter Wu, *Broadband Debate*]; Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 J. on Telecomm. & High Tech. L. 141, 145 (2003) [hereinafter Wu, *Network Neutrality*]; van Schewick, Dissertation, *supra* note 2, at 3-4, 102-03, 237-78, 342-49, 362-64.

critical for economic growth,[41] but also increases the Internet's potential to create value in the social, cultural, and political domains.[42] 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.[43] Third, network neutrality rules aim to preserve the Internet's ability to improve democratic discourse, facilitate political organization and action, and provide a decentralized environment for social, cultural, and political interaction in which anybody can participate.[44]

Network neutrality rules also have social costs: First, they limit the evolution of the network's core. Second, they limit network providers' ability to realize all potential efficiency gains or optimize the network in favor of the applications of the day. Third, they may reduce network providers' profits. Fourth, like any regulatory regime, they may create costs of regulation that burden providers, users, and society as a whole.[45]

Thus, the decision to adopt network neutrality rules is based on a trade-off.[46] In a way, the first two costs are the price of a system that can evolve and

---

41. *See, e.g.*, VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 356-57; Wu, *Broadband Debate*, *supra* note 40, at 80-81; van Schewick, Dissertation, *supra* note 2, at 346-49.

42. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 359-61.

43. *Id.* at 361-64; *see also infra* Box 3; *infra* note 60 and accompanying text.

44. *See, e.g.*, Remarks of Jack M. Balkin at FCC Workshop on Speech, Democratic Engagement, and the Open Internet, December 15, 2009, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 22, 2009) [hereinafter Balkin Remarks], *available at* http://apps.fcc.gov/ecfs/document/view?id =7020355385; VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 359-62, 364-65; Jack M. Balkin, *The Future of Free Expression in a Digital Age*, 36 PEPP. L. REV. 427, 436-38 (2009) [hereinafter Balkin, *Future of Free Expression*]; Bill D. Herman, *Opening Bottlenecks: On Behalf of Mandated Network Neutrality*, 59 FED. COMM. L.J. 103, 112-14 (2006); Marvin Ammori, *Net Neutrality and the 21st Century First Amendment*, BALKINIZATION (Dec. 10, 2009, 10:54 AM), http://balkin.blogspot.com/2009/12/net -neutrality-and-21st-century-first.html. On the Internet's social, cultural, and political potential, see generally YOCHAI BENKLER, THE WEALTH OF NETWORKS: HOW SOCIAL PRODUCTION TRANSFORMS MARKETS AND FREEDOM (2006).

45. *See, e.g.*, VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 365-68. These costs feature prominently, of course, in the arguments made by network neutrality opponents as well. *See, e.g.*, Gary S. Becker et al., *Net Neutrality and Consumer Welfare*, 6 J. COMPETITION L. & ECON. 497, 512-13, 516-19 (2010); Gerald R. Faulhaber, *Economics of Net Neutrality: A Review*, 3 COMM. & CONVERGENCE REV. 7, 22 (2011); Gerald R. Faulhaber & David J. Farber, *The Open Internet: A Customer-Centric Framework*, 4 INT'L J. COMM. 302, 313-15, 317-18, 321-24 (2010); Christopher S. Yoo, *Beyond Network Neutrality*, 19 HARV. J.L. & TECH. 1, 18-70 (2005); David Farber & Michael Katz, *Hold Off on Net Neutrality*, WASH. POST (Jan. 19, 2007), http://wapo.st/1AzOblI.

46. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 368-71; Frischmann & van Schewick, *supra* note 36, at 423-25. For a detailed analysis of the trade-off from the perspective of a network neutrality opponent, see Yoo, *supra* note 45, at 60-68, 70-76.

support new applications in the future.[47] 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.[48] As Tim Wu put it, it is like selling the painting to get a better frame.[49] Based on this reasoning, proponents of network neutrality resolve the trade-off in favor of the social benefits.[50]

Any network neutrality rule will impose these types of social costs, but different proposals for a nondiscrimination rule will support the goals of network neutrality regulation to varying degrees and will have different social costs. Most generally, policymakers should choose the rule that realizes the goals of network neutrality regulation and imposes the least social cost.[51]

In line with these considerations, a nondiscrimination rule (or any other network neutrality rule) should meet the following criteria[52]:

First, as I have explained elsewhere,[53] there are a number of factors that have allowed the Internet to foster application innovation, improve democratic

---

47. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 368-70.

48. *Id.* at 370-71.

49. Tim Wu, *Why You Should Care About Network Neutrality*, SLATE (May 1, 2006, 4:35 PM), http://www.slate.com/id/2140850.

50. *See, e.g.*, *Task Force Hearing*, *supra* note 36, at 56-57 (prepared statement of Timothy Wu, Professor of Law, Columbia Law School); VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 368-71; Frischmann & van Schewick, *supra* note 36, at 423-25.

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. *See, e.g.*, Comments of Cisco Systems, Inc. at 5-8, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Comments of Cisco Systems], *available at* http://fjallfoss.fcc.gov /ecfs/document/view?id=7020374147; Becker et al., *supra* note 45, at 506; Yoo, *supra* note 45, at 48-53.

51. This is not a strict optimization problem. The different types of social costs not only may be difficult to quantify exactly, but also may be incommensurable.

52. *See, e.g.*, Barbara van Schewick, Network Neutrality: What A Non-Discrimination Rule Should Look Like at 4-6, Attachment to Ex Parte Letter, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52, Framework for Broadband Internet Service, GN Docket No. 10-127 (Dec. 14, 2010) [hereinafter van Schewick, December 2010 Ex Parte Letter], *available at* http://apps.fcc.gov/ecfs/document /view?id=7020923837 (asking the FCC to use these criteria in evaluating alternative proposals for nondiscrimination rules in the context of the Open Internet proceeding); Barbara van Schewick, Network Neutrality: What a Non-Discrimination Rule Should Look Like at 1, Attachment to Notice of Ex Parte Conversation, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Aug. 2, 2010) [hereinafter van Schewick, August 2010 Attachment], *available at* http://fjallfoss.fcc.gov/ecfs /document/view?id=7020652518 (same).

discourse, facilitate political organization and action, and provide a more de-centralized environment for social, cultural, and political interaction in which anybody can participate. These factors need to be preserved to allow the Internet to continue to do so in the future,[54] and they 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.[55] They are the following[56]:

---

53. This paragraph is adapted from Barbara van Schewick, Notice of Ex Parte Meetings at 1-2, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 8, 2010) [hereinafter van Schewick, Notice of Ex Parte Meetings], *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=7020922658.

54. *See infra* note 56.

55. *See, e.g.*, van Schewick, Notice of Ex Parte Meetings, *supra* note 53, at 1-2 (arguing that 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"); Barbara van Schewick, Assistant Professor of Law, Stanford Law Sch., Opening Statement at the Federal Communications Commission's Workshop on Innovation, Investment and the Open Internet, Attachment to Notice of Ex Parte Meetings, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No 07-52 (Jan. 19, 2010) [hereinafter van Schewick, Innovation Opening Statement], *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id =7020382126 (asking the FCC to choose network neutrality rules that preserve application-blindness, user choice, innovation without permission, and low costs of application innovation).

56. The factors that have fostered application innovation in the past are described in detail in VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 12 tbl.I.2 (pointing to the parts of the book discussing these factors). For a short overview, see van Schewick, Innovation Opening Statement, *supra* note 55. For a brief discussion of the factors that are at the core of the Internet's political, social, and cultural potential, see Balkin Remarks, *supra* note 44; VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 359-65; and Yochai Benkler, *From Consumers to Users: Shifting the Deeper Structures of Regulation Toward Sustainable Commons and User Access*, 52 FED. COMM. L.J. 561, 565-68 (2000). 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 VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 61-75, 96-103; and van Schewick, Dissertation, *supra* note 2, at 81-109, 114-29. 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 Written Ex Parte of Professor Mark A. Lemley & Professor Lawrence Lessig at 1-2, 5-12, Application for Consent to the Transfer of Control of Licenses MediaOne Group, Inc. to AT&T Corp., CS Docket No. 99-251 (Nov. 10, 1999) [hereinafter Lemley & Lessig, Written Ex Parte], *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=6009850930 (discussing the issue in the context of the debate over open access to cable networks). For a similar discussion in the context of network neutrality, see *Hearing on the Future of the Internet*, *supra* note 33, at 52-57 (prepared statement of Lawrence Lessig, Professor of Law, Stanford Law School); *Hearing on Network Neutrality*, *supra* note 32, at 8-14 (prepared statement of Vinton G. Cerf, Vice President & Chief Internet Evangelist, Google Inc.); *Hearing on Network Neutrality*, *supra* note 32, at 54-59 (prepared statement of Lawrence Lessig, Professor of Law, Stanford Law School); *Government Role in Promoting the Future of the Telecommunica-*

- *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. Adding additional decisionmakers who need to endorse the idea or take action before an idea can be realized reduces the chances that innovative ideas can be realized.[57]

- *User Choice.* Users independently choose which applications they want to use without interference from network providers.[58] Letting users, not network providers, choose which applications will be successful is an important part of the mechanism that produces innovation under uncertainty.[59] 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 for them.[60] (See Box 3: The Importance of User Choice below.)

- *Application-Agnosticism.* The network is application-agnostic. While an application-agnostic network may have some information about the applications on the network, it does not make distinctions among data packets based on that information.[61] This ensures that network provid-

---

tions Industry and Broadband Deployment: Hearing Before the S. Comm. on Commerce, Sci. & Transp., 107th Cong. 33-34 (2002) (prepared statement of Lawrence Lessig, Professor of Law, Stanford Law School); Timothy Wu & Lawrence Lessig, Ex Parte Letter at 2-7, Appropriate Regulatory Treatment for Broadband Access to the Internet over Cable Facilities, CS Docket No. 02-52 (Aug. 22, 2003), *available at* http://apps.fcc.gov/ecfs/document/view?id=6514683885; LESSIG, *supra* note 40, at 34-46, 153-68, 246-49; Wu, *Broadband Debate*, *supra* note 40, at 145-51, 154-55, 170-72; Wu, *Network Neutrality*, *supra* note 40; van Schewick, Dissertation, *supra* note 2, at 3-4, 102-03, 237-349, 362-64; and note 2 above.

57. On innovation without permission in the original Internet, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 204, 211, 293. On the impact of innovation without permission on innovation, see *id.* at 345-48. *See also Hearing on Network Neutrality*, *supra* note 32, at 8-10 (prepared statement of Vinton G. Cerf, Vice President & Chief Internet Evangelist, Google Inc.); Balkin Remarks, *supra* note 44 (focusing on the social, cultural, and political implications).

58. *See Hearing on Network Neutrality*, *supra* note 32, at 8-9, 13 (prepared statement of Vinton G. Cerf, Vice President & Chief Internet Evangelist, Google Inc.); VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 144, 152-55, 293-95, 362-64.

59. *See* van Schewick, Innovation Opening Statement, *supra* note 55, at 6; *see also* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 349-51; *infra* note 60.

60. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 362-63; *see also Hearing on Network Neutrality*, *supra* note 32, at 8-9, 13 (prepared statement of Vinton G. Cerf, Vice President & Chief Internet Evangelist, Google Inc.). On the importance of user choice for the Internet's social, cultural, and political potential, see, for example, Balkin Remarks, *supra* note 44; and VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 359-65.

61. 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* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 72-75, 217-18; van Schewick, Dissertation, *supra* note 2, at 101-03; *supra*

ers cannot interfere with innovators' and users' choices, that they cannot distort competition among applications (or classes of applications), and that they cannot reduce application developers' profits through access fees.[62] We may call this "innovation without fear." (On the relationship between application-agnosticism and application-blindness, see Box 4: Application-Agnostic vs. 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 group of people to become innovators.[63] If there is uncertainty (for example, 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.[64] In the current Internet, there is uncertainty

---

note 2; *infra* Box 4; *see also, e.g.*, *Hearing on Network Neutrality*, *supra* note 32, at 8-10, 13 (prepared statement of Vinton G. Cerf, Vice President & Chief Internet Evangelist, Google Inc.); Lemley & Lessig, Written Ex Parte, *supra* note 56, at 7; Reed, *supra* note 2. For a short summary of the importance of application-blindness, see van Schewick, Innovation Opening Statement, *supra* note 55, at 3-4. For a detailed analysis, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 215-81, 286-95, 349-53, 355-65. 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 infra* Box 4. Thus, their economic, social, cultural, and political impact is the same. *See also* Balkin Remarks, *supra* note 44; VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 359-65 (focusing on the social, cultural, and political implications); Benkler, *supra* note 56, at 565-68.

62. 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 Article. On access fees, see, for example, van Schewick, Open Internet Opening Statement, *supra* note 29. *See also supra* Box 2; *supra* notes 29-30 and accompanying text.

63. For a short version of the argument, see van Schewick, Innovation Opening Statement, *supra* note 55, at 2-3, 5-6; and van Schewick, Open Internet Opening Statement, *supra* note 29, at 4-5. On the low cost of application innovation in the original Internet, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 138-48, 204-05, 289-90. On the impact of low-cost innovation on who can innovate, see *id.* at 204-13, 292-93. *See also* Balkin Remarks, *supra* note 44 (focusing on the social, cultural, and political implications); Benkler, *supra* note 56, at 565-68 (same).

64. For a short version of the argument, see van Schewick, Innovation Opening Statement, *supra* note 55, at 5-6; and van Schewick, Open Internet Opening Statement, *supra* note 29, at 4-5. For a detailed version, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 298-349.

and user needs are heterogeneous, so the conditions under which inno-
vator diversity increases the amount and quality of innovation are
met.[65]

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

Finally, the rule should make it easy to determine which behavior is and is
not allowed in order to provide certainty for industry participants. For applica-
tion providers, uncertainty over the level of protection provided by the rule re-
duces their incentives to innovate and their ability to get funding.[66] For net-
work providers, uncertainty over the legality of network management practices
or of different forms of Quality of Service may make it more difficult to man-
age their network and may limit the evolution of the network infrastructure.
Uncertainty over the regulatory regime may also reduce network providers' in-
centives to invest more generally.[67] Thus, certainty increases the social benefits
and reduces the social costs of a nondiscrimination rule.

In sum, policymakers should look for a rule that fosters application innova-
tion, protects user choice, and preserves the Internet's economic, social, cultur-
al, and political potential while avoiding unnecessary social costs. In particular,
the rule should preserve user choice, innovation without permission, applica-
tion-agnosticism, and low costs of application innovation. As this Article will
show, this framework not only can help guide the choice among alternative
proposals for nondiscrimination rules or other network neutrality rules, but also
can be used to interpret existing nondiscrimination rules or evaluate specific
forms of discriminatory conduct.

BOX 3
THE IMPORTANCE OF USER CHOICE[68]

Why the emphasis on user choice? First, user choice is fundamental if the
Internet is to create the maximum value to society. As a general-purpose tech-
nology, the Internet creates value by enabling users to do the things they want
or need to do. Users, not network providers, understand best which use of the

---

65. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 356.

66. *See infra* notes 236-37 and accompanying text.

67. *See infra* notes 233-35 and accompanying text.

68. This Box is adapted from Barbara van Schewick, Assistant Professor of Law, Stan-
ford Law Sch., Official Testimony at the Federal Communications Commission's Second
Public En Banc Hearing on Broadband Network Management Practices at 7-8, Broadband
Industry Practices, WC Docket No. 07-52 (Apr. 17, 2008) [hereinafter van Schewick, Offi-
cial Testimony], *available at* www.fcc.gov/broadband_network_management/041708
/vanschewick-written.pdf.

network is most valuable to them at any given time. Of course, in order for users to behave efficiently, they also need to bear (at least some of) the costs of their actions, something that 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, but not necessarily in the way that 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.

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.[69] 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 am I arguing for 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

---

69. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 72-74, 217-18; van Schewick, Dissertation, *supra* note 2, at 101-03; *supra* note 2; *see also, e.g.*, *Hearing on Network Neutrality*, *supra* note 32, at 8-10, 13 (prepared statement of Vinton G. Cerf, Vice President & Chief Internet Evangelist, Google Inc.); Reed, *supra* note 2.

use may enable network providers to predict or at least observe changes in the behavior of users, which may facilitate capacity planning.[70] Similarly, a clear overview of the applications using a network at a specific point in time may make it easier to detect security attacks.[71] At the same time, the potential harm to application innovation and user choice arises not 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, for instance, 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.[72]

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, but also threatens users' privacy.[73] Network neutrality rules address the first

---

70. Marjory S. Blumenthal & David D. Clark, *Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World*, 1 ACM TRANSACTIONS ON INTERNET TECH. 70, 77-78, 86 (2001); David D. Clark et al., *Tussle in Cyberspace: Defining Tomorrow's Internet*, 13 IEEE/ACM TRANSACTIONS ON NETWORKING 462, 472 (2005); Internet Eng'g Task Force, RFC 3639, Considerations on the Use of a Service Identifier in Packet Headers 4 (M. St. Johns & G. Huston eds., Oct. 2003), https://tools.ietf.org/html/rfc3639. This sentence and the next draw on VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 73.

71. Internet Eng'g Task Force, *supra* note 70, at 4.

72. 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, ARCHITECTURE AND INNOVATION, *supra* note 2, at 74-75.

73. *What Your Broadband Provider Knows About Your Web Use: Deep Packet Inspection and Communications Laws and Policies: Hearing Before the Subcomm. on Telecomms. & the Internet of the H. Comm. on Energy & Commerce*, 110th Cong. 61-62, 68-74, 77-84 (2008) (statement of David P. Reed, Adjunct Professor, Media Lab, Massachusetts Institute of Technology); Alissa Cooper, *Doing the DPI Dance: Assessing the Privacy Impact of Deep Packet Inspection*, *in* PRIVACY IN AMERICA: INTERDISCIPLINARY PERSPECTIVES 139, 146-53 (William Aspray & Philip Doty eds., 2011).

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.[74]

In its Open Internet Order, the FCC adopted a framework similar to the one described above. The Open Internet Rules set forth in this Order are explicitly based on the broader theoretical framework supported by most network neutrality proponents and used here.[75] Preserving user choice and innovation without permission is an explicit purpose of the Rules.[76] 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, the text of the Order explicitly specified that the FCC would 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' nondiscrimination rule.[77] (The same factors would be used to evaluate discriminatory or exclusionary conduct under the Rules' exception for reasonable network management.) Thus, the FCC explicitly specified that it would interpret key provisions of its rules—the nondiscrimination 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 nondiscrimination rules and specific discriminatory conduct throughout this Article: user choice, application-

---

74. 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 recreate 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 Paul Ohm, *The Rise and Fall of Invasive ISP Surveillance*, 2009 U. ILL. L. REV. 1417, 1489-96.

How privacy law should react to the erosion of application-blindness in today's Internet is beyond the scope of this Article. On this question, see, for example, Cooper, *supra* note 73, at 150-61; and Ohm, *supra*, at 1489-96.

75. 47 C.F.R. § 8.1 (2014); *Open Internet Order*, 25 FCC Rcd. 17,905, 17,908-15, 17,932-33, 17,949-50 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

76. 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*." (emphases added)).

77. See the discussion of the FCC's nondiscrimination rule in Part III below.

agnosticism, and innovation without permission.[78] The results of this Article indicate how the nondiscrimination rule and reasonable network 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.

II.   Proposals for Nondiscrimination 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 whether they should also ban discrimination. If they decide to ban discrimination, then they must also determine 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.[79]

This Part provides a comprehensive overview of existing proposals for nondiscrimination rules and evaluates them against the criteria developed in Part I to help policymakers choose among the available options. In the process, it explains how the different nondiscrimination rules affect network providers' ability to offer Quality of Service and which, if any, forms of Quality of Service a nondiscrimination rule should allow.

Part II.A defines the range of behaviors to which nondiscrimination rules apply. Nondiscrimination rules apply to any form of differential treatment that falls short of blocking. This includes the differential handling of data packets in the network, but also differential pricing practices directed towards subscribers.[80] Thus, nondiscrimination rules cover, for example, network provider practices that count only some applications, but not others, towards a subscriber's monthly bandwidth cap, or pricing plans that charge subscribers different Internet service fees for different applications or types of applications.

---

78. 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 Rules, including the nondiscrimination rule.

79. On Quality of Service, see notes 15-19 above. The following overview draws in part on Barbara van Schewick, Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like, at iv-xvii (2012), *available at* http:// cyberlaw.stanford.edu/files/publication/files/20120611-NetworkNeutrality_0.pdf, and van Schewick, *supra* note 20.

80. The nondiscrimination discussed in this Article do not address a network provider's pricing practices with respect to application providers that are not its Internet service customers. Whether an Internet service provider can charge application providers who are not its subscribers an "access fee" for access or for prioritized or otherwise enhanced access to its subscribers is governed by a network neutrality regime's rules about access fees. For a definition of the term "access fees," see note 62 above. A discussion of access fees is outside the scope of this Article. For a short overview of the debate, see Box 2 and notes 29-30 above.

A first set of proposals, discussed in Part II.B, takes an all-or-nothing approach to nondiscrimination. These approaches ban or allow all forms of differential treatment. While they provide a high degree of certainty and are easy to enforce, they do not accurately distinguish socially beneficial from socially harmful discrimination. As a result, they are either overinclusive or underinclusive and should be rejected.

A second set of proposals, discussed in Part II.C, recognizes that some forms of differential treatment will be socially harmful, while others will be socially beneficial, but assumes that it is impossible to distinguish among them in advance. Therefore, these proposals suggest adopting standards that specify criteria that will be used to judge specific discriminatory conduct in the future. Proposals in this category are afflicted with two problems: First, depending on the standard, they do not necessarily accurately separate socially beneficial from socially harmful conduct. For example, network neutrality proponents usually think of discriminatory conduct that favors an application over others as a distortion of competition and, therefore, as "anticompetitive." However, since antitrust law is based on a different theoretical framework that only considers a narrow range of economic harms, the term "anticompetitive" has a much narrower scope in antitrust law. As a result, a rule that bans behavior that is anticompetitive or violates an antitrust framework does not capture all instances of discrimination that threaten the values that network neutrality rules are designed to protect. Second, proposals in this category leave all decisions over the legality of specific discriminatory conduct to future adjudications. This creates considerable social costs. Case-by-case approaches fail to provide much-needed certainty to industry participants. They create high costs of regulation and tilt the playing field against those—end users, low-cost innovators, start-ups, nonprofits, independent artists, and members of underserved communities—who do not have the resources to engage in extended fights over the legality of specific discrimination in the future. They are also unlikely to lead to decisions that adequately protect the values network neutrality rules are intended to protect. In spite of these considerable social costs, the strategic interests of policymakers and of the big stakeholders on both sides of the network neutrality debate are aligned in favor of open-ended, case-by-case approaches. Thus, it is not surprising that many proposals in this category emerged from negotiations at the FCC or in Congress, or from direct negotiations between big stakeholders on opposite sides of the debate.

A final group of proposals would adopt more nuanced rules that specify in advance which differential treatment is and is not allowed. They are discussed in Part II.D. By adopting more nuanced criteria than all-or-nothing approaches, these proposals aim to more accurately distinguish between socially beneficial and socially harmful conduct (avoiding the problems of the all-or-nothing approaches), while doing so ex ante (avoiding the social costs of the standards-based approaches).

Of the three proposals in this category, only one, explored in Part II.D.2.b, accurately separates socially beneficial from socially harmful discrimination. It

would ban application-specific discrimination, but allow all application-agnostic discrimination. This is the rule policymakers should adopt. 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 the network to evolve; for example, it allows network providers to offer certain forms of user-controlled Quality of Service. As a bright-line rule, the rule provides certainty to market participants, keeps the cost of regulation low, and makes it feasible for users, start-ups, and nonprofits to navigate the process. Technically, it reinforces key architectural principles on which the Internet was based without locking in the original architecture of the Internet itself.

The two other proposals in this category—ban discrimination that is not disclosed and ban discrimination that does not treat like traffic alike—do not accurately distinguish socially beneficial from socially harmful discrimination and should be rejected.

Participants on both sides of the debate often assume that nondiscrimination rules that ban discriminatory conduct that is not disclosed will be sufficient to prevent blocking and discrimination if there is competition in the market for Internet services. As Part II.D.1 shows, this assumption is not correct. The market for Internet services is characterized by a number of factors—incomplete customer information, product differentiation in the markets 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, but only to some degree. They do not remove any of the other problems. As a result, they leave network providers with a substantial degree of market power over their customers that enables them to restrict some applications and content on their network without losing too many Internet service customers. Disclosure rules 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, disclosure rules are not a substitute for substantive rules against blocking or discrimination, even if there is competition in the market for Internet access services. This insight is particularly relevant for the debate over wireless network neutrality in the United States and for the network neutrality debates in Europe, Canada, and Australia.

Many network neutrality proponents support nondiscrimination rules that would 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. (This requirement is often called "like treatment.") Part II.D.2.a shows, however, that nondiscrimination rules that require like treatment do not adequately protect the values that network neutrality

rules are designed to protect. In particular, 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. This allows network providers to deliberately or inadvertently distort competition among applications or classes of applications and to interfere with user choice. Like treatment 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. 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 high costs of regulation.

A. *Scope of Nondiscrimination Rules*

Nondiscrimination 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.[81]

The scope of nondiscrimination 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, nondiscrimination rules apply to any form of differential treatment that may make some applications, classes of applications, or uses relatively more attractive to users than others. For example, Internet service providers can favor certain applications over others by not counting them towards users' monthly bandwidth caps or by charging a lower bandwidth-adjusted price for these applications. Therefore, these forms of differential treatment are subject to the nondiscrimination rules described in this Article, regardless of whether the packets associated with the favored applications receive the same technical treatment in the network as nonfavored applications.

An Internet service provider can make certain applications relatively more attractive by not counting (or "zero-rating") the traffic associated with these applications towards subscribers' monthly bandwidth cap. Consider an Internet service provider that zero-rates its own streaming video application, while the traffic of all other applications is counted towards subscribers' bandwidth cap. (For an example, see Box 5: Differential Counting of Traffic Towards the

---

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

Monthly Bandwidth Cap: Comcast's Xfinity TV App for the Xbox below.) For users who have not exhausted their monthly bandwidth allowance, watching a video that produces 2 gigabytes (GB) of traffic via an unaffiliated application brings those users 2 GB closer to exhausting their bandwidth cap. By contrast, watching the same video via the Internet service provider's application does not reduce the amount of bandwidth available to users before they reach the bandwidth cap. Users who have exhausted the monthly bandwidth allowance and watch the video using the unaffiliated application 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 watching the video via the affiliated application will not face any consequences. Thus, although the data packets associated with different streaming video applications receive the same technical treatment in the network,[82] the practice of counting only some streaming video applications towards the monthly bandwidth cap makes those applications relatively more attractive and is subject to the nondiscrimination rules discussed in this Article.

BOX 5

DIFFERENTIAL COUNTING OF TRAFFIC TOWARDS THE MONTHLY BANDWIDTH CAP: COMCAST'S XFINITY TV APP FOR THE XBOX

In the spring of 2012, Comcast, the largest provider of Internet service in the United States, 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's Xbox Live Gold subscription service. At the time, Comcast's Internet service had a 250 GB monthly bandwidth cap.[83] Traffic associated with the Xfinity TV app to the Xbox did 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 and Netflix) did.[84] As explained in the text, this type of differential counting of traffic towards the monthly bandwidth cap is generally subject to the nondiscrimination rules discussed in this Article, even if the data packets asso-

---

82. In the example of Comcast discussed in Box 5, it is unclear whether Comcast also prioritizes traffic associated with the Xfinity TV app for the Xbox over other traffic traveling to and from the Xbox in addition to zero-rating that traffic. *See infra* Box 5.

83. Cathy Avgiris, *Comcast to Replace Usage Cap with Improved Data Usage Management Approaches*, COMCAST (May 17, 2012), http://corporate.comcast.com/comcast-voices/comcast-to-replace-usage-cap-with-improved-data-usage-management-approaches. In May 2012, Comcast suspended enforcement of the 250 GB monthly bandwidth cap and started trialing several different data usage management approaches in selected markets. *Id.*

84. *FAQs: Xbox 360*, COMCAST, http://xbox.comcast.net/faqs.html (last visited Jan. 7, 2015).

ciated with the different streaming video applications receive the same technical treatment in the network.[85] Comcast claimed, however, that the Xfinity TV app was not provided over the public Internet and was therefore not subject to the FCC's Open Internet Rules.[86] Whether this assessment is correct is outside the scope of this Article.

It is unclear whether Comcast also prioritizes traffic associated with the Xfinity TV app for the Xbox over other traffic traveling to and from the Xbox. At the time, tests showed and Comcast admitted that Xfinity TV app traffic received special marking. In tests, the marked Xfinity TV app traffic consistently outperformed unmarked Netflix traffic to the Xbox during times of congestion, but Comcast claimed it was not prioritizing this traffic.[87]

Internet service providers may also favor specific applications or classes of applications over others by charging their subscribers different Internet service fees for different applications or types of applications.[88] 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, for instance to extract more of the value that Internet telephony users realize from the use of that application.[89] Other things being equal, this increases the costs of using Internet telephony relative to other applications, making Internet te-

---

85. For further discussion, see Box 15 and note 373 below.

86. Tony Werner, *The Facts About Xfinity TV and Xbox 360: Comcast Is Not Prioritizing*, COMCAST (May 15, 2012), http://corporate.comcast.com/comcast-voices/the-facts -about-xfinity-tv-and-xbox-360-comcast-is-not-prioritizing.

87. *See* Andrew Dugan, *An IP Engineer and Consumer View of Xfinity Traffic Prioritization*, BEYOND BANDWIDTH (May 17, 2012), http://blog.level3.com/2012/05/17/an-ip -engineer-and-consumer-view-of-xfinity-traffic-prioritization; Stacey Higginbotham, *He Said, She Said: Is Comcast Prioritizing Traffic or Not?*, GIGAOM (May 15, 2012, 1:17 PM PT), http://gigaom.com/2012/05/15/he-said-she-said-is-comcast-prioritizing-traffic-or-not; Werner, *supra* note 86.

88. Currently, Internet service providers in the United States generally 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, ARCHITECTURE AND INNOVATION, *supra* note 2, at 273-75, 277-78. *See also infra* note 425.

89. For a real-world example, two providers of deep packet inspection technology for mobile networks, Allot Communications and Openet, described a pricing scheme in which subscribers would pay two cents per megabyte for Facebook traffic, three dollars per month for Skype traffic, fifty cents per month for YouTube traffic, and nothing for traffic associated with the mobile provider's content offerings. *See* Allot Commc'ns & Openet, Managing the Unmanageable: Monetizing and Controlling OTT Applications, FierceMarks Webinar 7, Attachment to Free Press Ex Parte Letter, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 14, 2010), *available at* http://apps.fcc.gov/ecfs/document/view?id=7020923750.

lephony relatively less attractive.[90] Thus, Internet access plans that charge different bandwidth-adjusted prices for different applications would be subject to the nondiscrimination rules in this Article, even if the data packets associated with the different applications receive the same treatment in the network.

B. *All-or-Nothing Approaches*

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

The first approach would *allow all forms of discrimination*, which is equivalent to not adopting a nondiscrimination rule. This rule would not impose any constraints on Quality of Service. Proponents of this approach focus on the social benefits of allowing Quality of Service and other forms of differential treatment.

The second approach would *ban all forms of discrimination*. This approach would require network providers to treat each packet the same, which, by definition, would make it impossible to offer Quality of Service. Proponents of this approach emphasize the potential social costs of allowing Quality of Service and other forms of differential treatment. In particular, they are concerned that network providers may use Quality of Service as a tool to distort competition among competing applications by offering Quality of Service selectively to one of several competing applications. In addition, they fear that allowing network providers to offer Quality of Service and charge for it may reduce the quality of the baseline service and reduce network providers' incentives to increase the capacity of their networks.

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. However, these advantages come at a cost. Differential treatment and Quality of Service are not always good or always bad (for example, different forms of Quality of Service have different social benefits and social costs), but all-or-nothing approaches treat all forms of discrimination in the same way. As a result, banning all discrimination is overinclusive because that would ban socially beneficial forms of discrimination and would restrict the evolution of the network more than necessary to protect the values that network neutrality rules are designed to protect. Allowing all discrimination is underinclusive because that would allow socially harmful forms of discrimination and would effectively make the rule against blocking meaningless. Thus, both approaches should be rejected.

---

90. On the potential harms resulting from differentiating among classes of applications, see notes 366-406 and accompanying text below.

1. *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.[91] They would not adopt a nondiscrimination rule and instead would allow all discrimination. This approach would not restrict network providers' ability to offer Quality of Service in any way. The FCC'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.[92]

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.[93] They argue that without Quality of Service, 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.[94] 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

---

91. In the United States, see, for example, Comments of Cisco Systems, *supra* note 50, at 5-8; Comments of David Clark et al. at 21-22, 25, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Comments of Clark et al.], *available at* http://apps.fcc.gov/ecfs/document /view?id=7020373725; Comments of AT&T Inc. at iii, Broadband Industry Practices, WC Docket No. 07-52 (June 15, 2007) [hereinafter Comments of AT&T Inc.], *available at* http:// apps.fcc.gov/ecfs/document/view?id=6519529324; Comments of the National Cable & Telecommunications Ass'n at 36, Broadband Industry Practices, WC Docket No. 07-52 (June 15, 2007) [hereinafter Comments of the National Cable & Telecommunications Ass'n], *available at* http://apps.fcc.gov/ecfs/document/view?id=6519529328; Verizon Broadband Industry Practices Comments, *supra* note 34, at 41-42, 47-56; and Faulhaber & Farber, *supra* note 45, at 316-17, 336. In Europe, see, for example, Cave & Crocioni, *supra* note 22, at 1-3; and AT&T et al., *supra* note 22, at 1-4.

92. *Open Internet Order*, 25 FCC Rcd. 17,905, 17,956-61 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

93. *See* Comments of AT&T Inc., *supra* note 91, at 52-53.

94. *See* Becker et al., *supra* note 45, at 511-12, 518-19; Faulhaber, *supra* note 45, at 13; Robert E. Litan & Hal J. Singer, *Unintended Consequences of Net Neutrality Regulation*, 5 J. TELECOMM. & HIGH TECH. L. 533, 569-70 (2007). Allowing Internet service providers to offer Quality of Service may also allow them to differentiate themselves from their competitors. *See, e.g.*, Yoo, *supra* note 45, at 29-33 (discussing the benefits of allowing network providers to differentiate their Internet service offerings in general); *see also* Comments of AT&T Inc., *supra* note 91, at 71-73; Verizon Broadband Industry Practices Comments, *supra* note 34, at 57-58; Comments of TELUS Communications Co. at 20, Review of the Internet Traffic Management Practices of Internet Service Providers, CRTC 2008-19 (Feb. 23, 2009) (Can.) [hereinafter Comments of TELUS], *available at* http://www.crtc.gc.ca/public /partvii/2008/8646/c12_200815400/1029656.pdf.

exists) would be too difficult.[95] Since technology is evolving rapidly, regulators are likely to get it wrong.[96] Even if legislators or regulators succeed in identifying criteria that accurately distinguish between beneficial and harmful discrimination when the regulation is enacted, these criteria may not be accurate in the future. For these reasons, the argument concludes, 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—discrimination—that is less costly and potentially more effective.[97] Blocking and discrimination have the same effect: 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. 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 the poor performance of the affected application or website to design flaws and happily switch to the network provider's supposedly superior offering. But while blocking and discrimination have the same effect, the costs of discrimination are much lower. If the network provider blocks an application, users will notice and may switch to another Internet service provider.[98] 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.[99]

Based on these considerations, we would expect network providers to prefer discrimination over outright blocking. Indeed, their 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 (CRTC) showed, most of the larger Canadian Internet service providers throttled or otherwise interfered with peer-to-peer file-sharing applications, but did not block them completely.[100] Similarly, in 2009, British

---

95. *See* Comments of AT&T Inc., *supra* note 91, at 83-85.

96. *See* Comments of TELUS, *supra* note 94, at 4; Becker et al., *supra* note 45, at 509; Yoo, *supra* note 45, at 67.

97. The following paragraph draws on VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 260-61.

98. A user who notices that his application is being blocked will not necessarily switch providers, even if he would have preferred to use the blocked application over alternative applications that are not blocked. *See infra* Part II.D.1.

99. Requiring network providers to disclose whether they interfere with applications and content will not fully solve this problem. *See infra* Part II.D.1.

100. CHRISTOPHER PARSONS, SUMMARY OF JANUARY 13, 2009 CRTC FILINGS BY MAJOR ISPS IN RESPONSE TO INTERROGATORY PN 2008-19 WITH FEBRUARY 9, 2009 UPDATES 23-31

telecommunications provider BT restricted the bandwidth available to the BBC iPlayer and other streaming video applications to 896 kilobits per second in BT's "Up to 8 Mbps Option 1" broadband service; it did not block these applications completely.[101] 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 connection.[102] RCN, a competitive cable provider in the United States, used the same technology from 2005 to 2009.[103] As white papers produced by Comcast's equipment vendor, Sandvine, showed, this method of interference was deliberately chosen to prevent customers from noticing it.[104] 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.[105]

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

---

(2009), *available at* http://www.christopher-parsons.com/PublicUpload/Summary _of_January_13_2009_ISP_filings_with_February_9_2009_Updates_version_1.0(for_web) .pdf.

101. *See* Rory Cellan-Jones, *iPlayer: BBC v BT*, BBC NEWS DOT.LIFE (June 2, 2009, 9:20 AM GMT), http://bbc.co.uk/blogs/technology/2009/06/iplayerbbc_v_bt.html.

102. *See* Comcast Corp. Description of Current Network Management Practices, Letter from Kathryn A. Zachem, Vice President, Regulatory Affairs, Comcast Corp., to Marlene Dortch, Sec'y, FCC, Attachment A, Formal Complaint of Free Press & Public Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, No. EB-08-IH-1518, Broadband Industry Practices, WC Docket No. 07-52 (Sept. 19, 2008) [hereinafter Comcast Description], *available at* http://apps.fcc.gov/ecfs/document/view?id=6520172537.

103. Like Comcast, RCN limited the number of simultaneous, unidirectional uploads and prevented additional uploads from occurring when the threshold had been reached; both also used the same Sandvine equipment (Sandvine PTS 8210). However, Comcast deployed the Sandvine Policy Switch out-of-line, while RCN deployed it in-line. *See id.* at 4, 5; *see also* RCN Corp., Ex Parte Notice at 1-4, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (May 7, 2010) [hereinafter RCN Letter], *available at* http://apps.fcc.gov/ecfs/document/view?id=7020450131.

104. *See* SANDVINE INC., MEETING THE CHALLENGE OF TODAY'S EVASIVE P2P TRAFFIC: SERVICE PROVIDER STRATEGIES FOR MANAGING P2P FILESHARING 14 (2004), *available at* http://www.larryblakeley.com/Articles/p2p/Evasive_P2P_Traffic.pdf.

105. *See* Eric Hellweg, *The Kazaa Conundrum*, CNN MONEY (Sept. 10, 2003, 1:15 PM EDT), http://money.cnn.com/2003/09/10/technology/techinvestor/hellweg/index.htm; Thomas Mennecke, *DSL Broadband Providers Perform Balancing Act*, SLYCK NEWS (Nov. 1, 2005), http://www.slyck.com/news.php?story=973; *see also* Comments of the National Cable & Telecommunications Ass'n, *supra* note 91, at 31 ("[C]able operators will not go down the path of blocking access to video or P2P services. Blocking such services would be a recipe for . . . massive dissatisfaction among consumers, which would lead to loss of customers to our competitors."); SANDVINE INC., *supra* note 104, at 5-6.

costs. Some are socially beneficial, and 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.

2.   *Ban all discrimination*

By contrast, some participants in the debate would ban all discrimination, requiring network providers to treat every packet the same.[106] The FCC's draft nondiscrimination rule in the Open Internet proceeding is an example of this type of approach.[107] A rule that required 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.[108]

---

106. *See* Crawford, *supra* note 32, at 403-04; *see also* Brett M. Frischmann, Notice of Ex Parte Conversation, Preserving the Open Internet, GN Docket No. 09-191 (Aug. 9, 2010), *available at* http://apps.fcc.gov/ecfs/document/view?id=7020664240; BRETT M. FRISCHMANN, INFRASTRUCTURE: THE SOCIAL VALUE OF SHARED RESOURCES 348-55 (2012) [hereinafter FRISCHMANN, INFRASTRUCTURE]. While Frischmann's proposed nondiscrimination 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. FRISCHMANN, INFRASTRUCTURE, *supra*, at 348-55; *see also* Frischmann, *supra* note 32, at 1011-12.

107. Preserving the Open Internet, 24 FCC Rcd. 13,064, 13,104 (proposed Oct. 22, 2009) ("Subject to reasonable network management, a provider of broadband Internet access service must treat lawful content, applications, and services in a nondiscriminatory manner." (italics omitted)); *id.* at 13,104-06 (noting that the proposed nondiscrimination 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" (emphasis omitted)). This rule would have been subject to an exception for reasonable network management. *Id.* at 13,113-15. This nondiscrimination rule was supported by, for example, Free Press, Free Press Open Internet Comments, *supra* note 3, at 74-75; the Open Internet Coalition, Open Internet Coalition Comments, *supra* note 36, at 15-17; and public interest commenters, Public Interest Comments, *supra* note 36, at v. These commenters would have coupled the nondiscrimination 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 (1) serve a legitimate purpose related to the technical management of the network and (2) be narrowly tailored to address that purpose. *See, e.g.*, Free Press Open Internet Comments, *supra* note 3, at 78-79, 82-104; Open Internet Coalition Comments, *supra* note 36, at 41-50; Public Interest Comments, *supra* note 36, at 35-44.

108. Nondiscrimination rules usually have an exception for reasonable network management. Thus, even under the strict nondiscrimination rule described in the text, network

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 applications, but not to other, competing applications, or may sell Quality of Service exclusively to one of several competing applications.[109] 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.[110] Moreover, selling Quality of Service allows network providers to profit from bandwidth scarcity, which reduces their incentives to increase the capacity of their networks.[111] 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.[112]

---

providers may still be able to provide some or all forms of Quality of Service, provided that 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.

109. Free Press Open Internet Comments, *supra* note 3, at 21-23; Public Interest Comments, *supra* note 36, at 48, 51.

110. *See, e.g.*, Nicholas Economides, *Why Imposing New Tolls on Third-Party Content and Applications Threatens Innovation and Will Not Improve Broadband Providers' Investment*, *in* NET NEUTRALITY: CONTRIBUTIONS TO THE DEBATE 87, 94 (Jorge Pérez Martínez ed., 2010).

111. Center for Democracy & Technology Comments, *supra* note 36, at 28-29; Free Press Open Internet Comments, *supra* note 3, at 22; Open Internet Coalition Comments, *supra* note 36, at 46; Public Interest Comments, *supra* note 36, at 45; Economides, *supra* note 110, at 94, 99-100.

112. First, the nondiscrimination 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 infra* Part II.D.2.b.i.A; *infra* note 483 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* van Schewick, Background Paper, *supra* note 29, at 10-11; *infra* note 482 and accompanying text. 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* van Schewick, Innovation Opening Statement, *supra* note 55, at 4-6; van Schewick, Open Internet Opening Statement, *supra* note 29, at 3-5; van Schewick, Background Paper, *supra* note 29, at 10-12; *infra* note 477 and accompanying text.

Supporters of banning Quality of Service also question whether Quality of Service is needed at all.[113] If there is no need for Quality of Service, then banning it creates limited social costs.[114] 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.[115] 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.[116] Video telephony applications like Skype or Google Video Chat function over today's broadband connections.[117] The success of real-time applications on today's best-effort Internet is due to two reasons: First, many regions currently seem to have sufficient network capacity to prevent the lack of Quality of Service from becoming a problem.[118] 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.[119] Second, 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.[120] 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.[121]

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,

113. *See* Open Internet Coalition Comments, *supra* note 36, at 33-35; Public Interest Comments, *supra* note 36, at 49-50.

114. For a similar interpretation, see Comments of Clark et al., *supra* note 91, at 10.

115. *See* Open Internet Coalition Comments, *supra* note 36, at 33-35; Public Interest Comments, *supra* note 36, at 49-50.

116. PETERSON & DAVIE, *supra* note 15, at 531.

117. For example, Skype video requires a high-speed broadband connection of at least 512 kilobits per second (kbps) download and 128 kbps upload. For best quality, Skype recommends "a high-speed broadband connection of 4Mbps down / 512kbps up." *Skype for Windows Desktop: System Requirements*, SKYPE, http://www.skype.com/intl/en-us/get-skype/on-your-computer/windows (last visited Jan. 7, 2015).

118. Bruce Davie, *Deployment Experience with Differentiated Services*, 2003 PROC. ACM SIGCOMM 2003 WORKSHOPS 131, 134.

119. *See* KUROSE & ROSS, *supra* note 15, at 629-30. For a more detailed analysis of the relationship among capacity, congestion, and Quality of Service, see the discussion in the following paragraphs.

120. *See id.* at 616-29.

121. *See, e.g.*, Open Internet Coalition Comments, *supra* note 36, at 43-46; FRISCHMANN, INFRASTRUCTURE, *supra* note 106, at 353-55.

creating congestion.[122] (See Box 6: The Relationship Between Congestion, Delay, Jitter, and Loss below.) 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.[123] 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[124] or may provide a performance that is more constant and predictable than best-effort service.[125] These services, however, can improve on the performance of best-effort service only if there is congestion.[126] If there is no congestion (i.e., if all queues are empty), congestion-related loss and queuing delay will constantly be zero, jitter will be low for all packets, and data flows

---

122. The definition of congestion used throughout this Part 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 7 below.

123. 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 factors, loss, and delay. Combinations of two of these three variables determine the value of the third. In particular, "[f]or 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." NEIL DAVIES ET AL., AN OPERATIONAL MODEL TO CONTROL LOSS AND DELAY OF TRAFFIC AT A NETWORK SWITCH 3 (1999), *available at* http://www.cs.bris.ac.uk/Publications/Papers/1000387.pdf. The variability of delay in turn determines jitter.

124. 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. *See* Geoff Huston, *The QoS Emperor's Wardrobe*, ISP COLUMN 2 (June 2012), http://www.potaroo.net/ispcol/2012-06/noqos.pdf.

125. 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 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. *Id.*

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* CISCO SYS., ENTERPRISE QOS SOLUTION REFERENCE NETWORK DESIGN GUIDE 1-22 (2005), *available at* http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.pdf; Stanislav Shalunov & Benjamin Teitelbaum, Internet2, QBone Scavenger Service (QBSS) Definition (Mar. 16, 2001) (on file with author).

126. How the performance of the service compares with best-effort service in the absence of congestion depends on the type of service. *See supra* notes 124-25 and accompanying text.

will experience the maximum throughput and minimum end-to-end delay that is possible on their path.[127] No class of service can improve on that. Thus, Quality of Service is only useful if there is at least some congestion.

BOX 6

THE RELATIONSHIP BETWEEN CONGESTION, DELAY, JITTER, AND LOSS

Throughout this Part, "congestion" denotes the building up of a queue for an outgoing link at a router, which may increase delay, jitter, or packet loss.[128] (This definition differs from the definition of congestion that is often used by network providers. See Box 7: 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.[129] If packets destined for a specific outgoing link arrive faster than they can be 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.[130] As the

---

127. 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, *supra* note 15, at 36-40 (describing the different types of delay contributing to a packet's total end-to-end delay).

128. The definition of congestion used throughout this Part is derived from the definition of congestion in queuing theory. *See infra* Box 7.

129. On the following, see Testimony of Doctor David Reed ¶¶ 6-18, Initial Comments of Campaign for Democratic Media, Attachment B, Part 1, Review of the Internet Traffic Management Practices of Internet Service Providers, CRTC 2008-19 (Feb. 23, 2009) (Can.) [hereinafter Reed CRTC Testimony], *available at* http://www.crtc.gc.ca/public/partvii/2008 /8646/c12_200815400/1029987.zip. *See generally* KUROSE & ROSS, *supra* note 15, at 337-40, 653-54; PETERSON & DAVIE, *supra* note 15, at 16-17, 479-80, 492-93. The text describes the scheduling and drop policy—first in, first out (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 discard packets from its queue to make space for the arriving packet. *See* PETERSON & DAVIE, *supra* note 15, at 492-93.

130. *See, e.g.*, PETERSON & DAVIE, *supra* note 15, at 16-17; *BufferBloat: What's Wrong With the Internet?*, COMM. ACM, Feb. 2012, at 40, 43. For a more detailed description of potential reasons for congestion, see Box 8 below.

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 (or, in other words, how much congestion the packet encounters along its path), 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.[131] 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.[132] 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.

BOX 7

DEFINITIONS OF CONGESTION AND BENEFITS FROM QUALITY OF SERVICE

Throughout this Part, "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 6: The Relationship Between Congestion, Delay, Jitter, and Loss.) This definition is derived from the definition of congestion used in queuing theory.[133] As explained in the text, Quality of Service only provides an improvement over best-effort service if this type of congestion exists.

---

131. 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. *See* KUROSE & ROSS, *supra* note 15, at 37-40; PETERSON & DAVIE, *supra* note 15, at 46-47.

132. PETERSON & DAVIE, *supra* note 15, at 54-55.

133. *See* Steven Bauer et al., The Evolution of Internet Congestion 10 (Aug. 15, 2009) (unpublished manuscript), *available at* http://people.csail.mit.edu/wlehr/Lehr-Papers_files /Bauer_Clark_Lehr_2009.pdf ("According to [the queuing theory] definition, as soon as a queue starts to build traffic congestion is occurring."); *see also* Reed CRTC Testimony, *su-*

> 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.[134]
>
> While Quality of Service is useless in a network that never experiences congestion under the definition used throughout this Part, it 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.[135] Thus, a network that is not congested under the definition used by network providers may experience congestion under the definition used throughout this Part 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 Part'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.[136] In such a network, users may find Quality of Service very valuable and may be very willing to pay for it.[137]

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 delay. But even a network with low average utiliza-

---

pra note 129, ¶¶ 7, 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.").

134. *See* Reed CRTC Testimony, *supra* note 129, ¶¶ 24-25; Bauer et al., *supra* note 133, at 10-11. A network can be congested under the queuing theory definition but not under the network provider's definition and vice versa. Bauer et al., *supra* note 133, at 11. For a detailed analysis and comparison of different definitions of congestion, see *id.* at 8-13.

135. *See infra* Box 8; *infra* note 138.

136. While offering different types 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 a network that offers different types of service needs a certain amount of capacity to provide satisfactory performance. *See, e.g.*, KUROSE & ROSS, *supra* note 15, at 664-65; PETERSON & DAVIE, *supra* note 15, at 553.

137. 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 supra* notes 110-12 and accompanying text.

tion will experience occasional congestion.[138] 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, the performance of that application will suffer.[139] (See Box 8: Causes of Congestion in a Network with Low Average Utilization.)

BOX 8

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.[140] 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

---

138. This insight is well established in the literature. *See, e.g.*, Comments of Clark et al., *supra* note 91, at 10; Brian E. Carpenter & Kathleen Nichols, *Differentiated Services in the Internet*, 90 PROC. IEEE 1479, 1482-83 (2002) ("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 traffic is modest."); Bauer et al., *supra* note 133, at 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 [of] the Internet varies across the network, congestion events may arise commonly even in a network that may be considered to be generally 'over-provisioned.'").

139. The ability of an application to compensate for increases in 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 videoconferencing 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, *supra* note 15, at 532-34; Ilpo Järvinen et al., Impact of TCP on Interactive Real-Time Communication 1-2 (July 28, 2012) (unpublished manuscript), *available at* http://www.tschofenig.priv.at/cc-workshop/irtf_iab-ccirtcpaper9.pdf.

140. The relationships described in the text are explained in more detail in, for example, Reed CRTC Testimony, *supra* note 129, ¶¶ 16-18; KUROSE & ROSS, *supra* note 15, at 40-42; and PETERSON & DAVIE, *supra* note 15, at 54, 54.

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 rapid building up and emptying of queues not only increases the delay experienced by other applications that are transferring data over the same link at the same time, but also increases jitter. The increase in jitter and delay harms applications such as interactive voice and video applications or online gaming applications that need low jitter or delay.[141] Recent changes to transport protocols[142] and operating systems[143] have increased the amount of data a single Transmission Control Protocol (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.[144]

Applications that upload or download a lot of data using TCP (e.g., for uploading a video to YouTube, sending or receiving e-mails 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

---

141. *See* Järvinen et al., *supra* note 139 (discussing the problem of increased jitter and delay and presenting the results of an experiment that demonstrates the problem).

142. For example, a recent IETF experimental standard proposed an increase in the permitted upper bound for TCP's initial window (IW) to ten segments depending on the maximum segment size. J. Chu et al., Internet Eng'g Task Force, RFC 6928, Increasing TCP's Initial Window (Apr. 2013), https://tools.ietf.org/html/rfc6928. These changes allow each new TCP connection "to send as much as 2.5 times as much data as in the past." Jim Gettys, Internet Eng'g Task Force, Internet-Draft, IW10 Considered Harmful (Aug. 26, 2011), https://tools.ietf.org/html/draft-gettys-iw10-considered-harmful-00. For an analysis of these changes' impact on other applications that are transferring data at the same time, see Järvinen et al., *supra* note 139; Chu et al., *supra*, at 5, 10, 12-15, 21-23; and Gettys, *supra*.

143. For example, Windows XP did not implement TCP window scaling, and therefore it limited the number of packets it sent before it received an acknowledgment. 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 BufferBloat: What's Wrong With the Internet?*, *supra* note 130, at 44-45; Jim Gettys & Kathleen Nichols, *Bufferbloat: Dark Buffers in the Internet*, 55 COMM. ACM 57, 64 (2012). TCP window scaling is a TCP option that makes it possible to increase the size of TCP's receive window beyond 65,535 bytes, the maximum size of the receive window under normal TCP. V. Jacobson et al., Internet Eng'g Task Force, RFC 1323, TCP Extensions for High Performance (May 1992), https://www.ietf.org/rfc/rfc1323.txt.

144. 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* Järvinen et al., *supra* note 139, at 1; Gettys, *supra* note 142, at 3-4.

duration of the flow, which increases delay for other applications trying to transfer data at the same time.[145]

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.[146]

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 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 breakups 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.[147] 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, allowing 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 guaran-

---

145. *See* Kathleen Nichols & Van Jacobson, *Controlling Queue Delay*, ACM QUEUE, May 2012, at 1, 4-5; Järvinen et al., *supra* note 139, at 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 generally BufferBloat: What's Wrong With the Internet?*, *supra* note 130; Gettys & Nichols, *supra* note 143.

146. *See* Comments of Clark et al., *supra* note 91, at 10; Bauer et al., *supra* note 133, at 16.

147. 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 nondiscrimination rule should allow is discussed below in notes 430-36 and accompanying text (discussing forms of Quality of Service that treat like traffic alike) and notes 484-85 and accompanying text (discussing certain forms of user-controlled Quality of Service).

teed throughput, jitter, or delay.[148] 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.[149] 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.[150] 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.[151] 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, and search engines.[152] Thus, just because 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 innovators is allowed to innovate under the right conditions,[153] the innovators will find ways to use the Internet's functionality that those who originally designed that functionality had not necessarily thought of,

---

148. It is not possible to construct a transport-layer service that guarantees delay (or bandwidth) if, as in the current best-effort Internet, the Internet layer does not guarantee delay (or bandwidth). KUROSE & ROSS, *supra* note 15, at 97, 201; VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 142 box 4.3.

149. *See* Litan & Singer, *supra* note 94, at 569-70.

150. This argument has come up repeatedly in personal discussions with proponents of a ban.

151. *See, e.g.*, Nathan Rosenberg, *Uncertainty and Technological Change, in* THE MOSAIC OF ECONOMIC GROWTH 334, 334 (Ralph Landau et al. eds., 1996). The following sentences draw on VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 301-02.

152. For a more detailed discussion of these examples, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 301-04.

153. These conditions include the factors described above: innovation without permission, innovation without fear (or application-agnosticism), user choice, and low costs of application innovation. *See supra* notes 56-65 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 nondiscrimination rule should allow, see notes 407-40 and accompanying text (discussing forms of Quality of Service that treat like traffic alike) and notes 482-86 and accompanying text (discussing certain user-controlled forms of Quality of Service) below.

and at least some of the resulting applications or uses will create significant social value.[154]

Finally, in situations in which 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.[155]

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."[156] This solution is called "overprovisioning."[157] 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 overprovisioning residential access networks, or, as they described it, providing the "overabundance of bandwidth . . . [that] ensure[s] that the odds of network congestion are minimized," would require offering a 1 gigabit per second connection to residential users (where 1 gigabit per second equals 1000 megabits per second (Mbps)).[158] Since then, the demands and capabilities of end devices and applications have evolved rapidly, so the capacity required to overprovision access networks today will most likely be higher. For example, a single TCP connection on a personal computer can send data at a rate of hundreds of

---

154. For a more detailed discussion of this argument based on economic theory and the history of specific Internet applications, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, ch. 8, at 297-353.

155. See the discussion of application-agnostic network management coupled with user-controlled differentiation in Part II.D.2.b.i.B below.

156. Reed CRTC Testimony, *supra* note 129, ¶ 23.

157. As Bauer et al. explain, a common approach to provisioning, which provisions for expected peak demand over some time period, may often result in networks that are overprovisioned 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 off-peak 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.

Bauer et al., *supra* note 133, at 6.

158. STEVEN C. CORBATÓ & BEN TEITELBAUM, INTERNET2 AND QUALITY OF SERVICE: RESEARCH, EXPERIENCE, AND CONCLUSIONS 2-3 (2006), *available at* http://www.educause .edu/ir/library/pdf/CSD4577.pdf; *see also Reconsidering Our Communications Laws: Ensuring Competition and Innovation: Hearing Before the S. Comm. on the Judiciary*, 109th Cong. 219 (2006) [hereinafter *Hearing on Communications Laws*] (prepared statement of Jeff C. Kuhns, Senior Director, Consulting and Support Services, Pennsylvania State University); *Hearing on Network Neutrality*, *supra* note 32, at 65-66 (prepared statement of Gary R. Bachula, Vice President, Internet2).

megabits per second, so a single user could easily create peak rates of more than a gigabit per second by opening several TCP connections simultaneously.[159] Moreover, TCP is designed to 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.[160] Finally, even in an overprovisioned 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 create congestion as well.[161] Thus, while overprovisioning will further reduce the probability of congestion, it cannot eliminate it.[162] Due to the low likelihood of congestion, a network that is truly overprovisioned 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.[163]

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

---

159. Comments of Clark et al., *supra* note 91, at 10.

160. For a detailed explanation of this point, see *id.* ("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."); Bauer et al., *supra* note 133, at 16.

161. *See supra* Box 8.

162. For a number of reasons,

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 [of] the Internet varies across the network, congestion events may arise commonly even in a network that may be considered to be generally "over-provisioned."

Bauer et al., *supra* note 133, at 6; *see also* Ben Teitelbaum & Stanislav Shalunov, Internet2, Why Premium IP Service Has Not Deployed (and Probably Never Will) (Jan. 9, 2006) (on file with author) ("Although well-provisioned networks deliver very good typical performance, they will, in general, deliver unpredictable service and, in the worst case, no service.").

163. Benjamin Teitelbaum & Stanislav Shalunov, *What QoS Research Hasn't Understood About Risk*, 2003 PROC. ACM SIGCOMM 2003 WORKSHOPS 148, 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). Teitelbaum and Shalunov also argue that "premium service" would be valuable even in an overprovisioned 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, *supra* note 162.

capacity does not necessarily prevent congestion, and Quality of Service may therefore be useful in networks with more capacity as well. In networks that have low average utilization, but are not overprovisioned, 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.[164] Quality of Service may also enable new applications that we have not yet thought of 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 overprovisioning, Quality of Service may provide benefits even in overprovisioned networks by allowing users to 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 overprovisioning 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 anyway.[165] 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 overprovisioning is more efficient.[166] After developing and successfully testing Quality of Service technology in the research network Internet2 for several years, Internet2 researchers suspended the effort indefinitely.[167] While

---

164. Overprovisioning requires considerably more capacity than ensuring low average utilization, so a lot of networks may belong to this category.

165. *See, e.g.*, Open Internet Coalition Comments, *supra* note 36, at 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").

166. *See* KUROSE & ROSS, *supra* note 15, at 602-04; Gregory Bell, *Failure to Thrive: QoS and the Culture of Operational Networking*, 2003 PROC. ACM SIGCOMM 2003 WORKSHOPS 115, 116-19 (discussing the trade-off in the context of enterprise networks); Davie, *supra* note 118, at 134-35. For proponents of overprovisioning, see ANDREW ODLYZKO, THE ECONOMICS OF THE INTERNET: UTILITY, UTILIZATION, PRICING, AND QUALITY OF SERVICE (Ctr. for Discrete Mathematics & Theoretical Computer Sci., DIMACS Technical Report 99-08, 1999), *available at* http://citeseerx.ist.psu.edu/viewdoc/download ?doi=10.1.1.45.7943&rep=rep1&type=pdf; Andrew Odlyzko, *The Current State and Likely Evolution of the Internet*, 3 GLOBAL TELECOMM. CONF.: GLOBECOM '99, at 1869 (1999); Dan Bricklin, *Why We Don't Need QOS: Trains, Cars, and Internet Quality of Service*, DAN BRICKLIN'S WEBSITE (Aug. 2, 2003), http://www.bricklin.com/qos.htm; Huston, *supra* note 124; and sources cited in note 169 below.

167. Teitelbaum & Shalunov, *supra* note 162.

they acknowledged that being able to protect important applications against the risk of congestion is valuable even in an overprovisioned network, they concluded that "the costs . . . are too high relative to the perceived benefits" and that overprovisioning was the more efficient solution.[168] 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.[169]

While introducing Quality of Service creates costs, overprovisioning—which requires considerably more capacity than that needed to ensure 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 overprovision than access networks because they can take advantage of statistical aggregation. Overprovisioning research networks whose users are already attached to high-speed campus networks may be less costly than overprovisioning 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 multiprovider networks.[170] Today, many corporate intranets use Quality of Service; large Internet service providers give business customers the

---

168. *Id.* ("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 DiffServ 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." *Id.*

169. *Hearing on Communications Laws*, *supra* note 158, at 218-19 (prepared statement of Jeff C. Kuhns, Senior Director, Consulting and Support Services, Pennsylvania State University); *Hearing on Network Neutrality*, *supra* note 32, at 65-68 (prepared statement of Gary R. Bachula, Vice President, Internet2). The Internet2 experience is often cited by proponents of a ban on Quality of Service. *See, e.g.*, Open Internet Coalition Comments, *supra* note 36, at 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").

170. 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 multiprovider 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, *supra* note 118, at 134; Huston, *supra* note 124.

option of buying different classes of service.[171] Thus, whether overprovisioning 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 overprovisioning and Quality of Service is an important one that is 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.

Rather, in the context of the network neutrality debate, the only relevant question is whether banning Quality of Service is 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.[172] 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.[173] If regulators adopt nondiscrimination rules that allow certain forms of Quality of Service, they do not pick winners and losers in this debate. Such nondiscrimination 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 overprovisioning 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 allowing Quality of Service may indeed harm 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 than an all-or-nothing approach is needed. While the 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

---

171. On enterprise deployments, see PETERSON & DAVIE, *supra* note 15, at 554; and Davie, *supra* note 118, at 133-34. On Quality of Service offerings for business customers, see, for example, FALK VON BORNSTAEDT, DEUTSCHE TELEKOM GRP., QUALITY-OF-SERVICE (QOS) FOR IP-TRANSIT / PEERING 1, 3 (2012), *available at* https://bscw.ict-etics.eu/pub/bscw .cgi/d44579/DT%20QoS%20White%20Paper.pdf; and Verizon, Verizon Private IP 25, 29 (2006) (on file with author) [hereinafter Verizon Presentation].

172. This is an example of the general principle, described in Part I above, that while network neutrality rules may require imposing some restrictions on innovation in the network in order to allow the Internet to continue to foster application innovation, preserve user choice, and foster democratic discourse, they should not restrict the evolution of the network more than necessary to realize these goals. *See supra* notes 46-50 and accompanying text.

173. Of course, the constraints imposed by a nondiscrimination rule that allows all or some forms of Quality of Service will influence network providers' private costs and benefits of overprovisioning 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.

high average utilization. In networks that have low average utilization without being overprovisioned, Quality of Service may allow users to improve the performance of existing applications, enable new applications that benefit from the availability of different classes of service, and enable users whose bandwidth is limited to use that bandwidth more efficiently. Ensuring low average utilization requires considerably less capacity than overprovisioning, so many networks may belong to the category just described. In an overprovisioned 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.

C. *Case-by-Case Approaches*

A second set of approaches would 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 overinclusive. 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 the motivations for and impact of the practice, are known.[174] To support their proposals, they point to the example of antitrust law, which evaluates behavior that may be anticompetitive or procompetitive depending on the circumstances after the fact on a case-by-case basis.[175]

Approaches in this group differ along two dimensions[176]: the degree to which they prescribe the standard that regulators should use to assess specific

---

174. *See, e.g.*, Robert Hahn et al., *Addressing the Next Wave of Internet Regulation: Toward a Workable Principle for Nondiscrimination*, 4 REG. & GOVERNANCE 365, 368 (2010); Philip J. Weiser, *Toward a Next Generation Regulatory Strategy*, 35 LOY. U. CHI. L.J. 41, 75-76 (2003); Christopher S. Yoo, *What Can Antitrust Contribute to the Network Neutrality Debate?*, 1 INT'L J. COMM. 493, 515-17 (2007).

175. *See, e.g.*, Weiser, *supra* note 174, at 75-76; Yoo, *supra* note 174, at 515-16.

176. Proposals also differ on which institutional actor should perform the adjudication in the United States: the Federal Trade Commission (FTC), *see, e.g.*, Jonathan E. Nuechterlein, *Antitrust Oversight of an Antitrust Dispute: An Institutional Perspective on the Net Neutrality Debate*, 7 J. ON TELECOMM. & HIGH TECH. L. 19, 57-65 (2009), or the FCC, *see, e.g.*, Hahn et al., *supra* note 174, at 366, 368 (listing reasons for enforcement by the FCC without taking a side in the debate); Weiser, *supra* note 174, at 75, 77-78. This question is outside the scope of this Article. The FTC is only a plausible option for those who base network neutrality regulation on an antitrust framework. If network neutrality regulation is

discriminatory behavior, and the extent to which the approaches are able to capture the instances of discrimination that threaten the values that network neutrality rules are designed to protect. Taken together, these two characteristics determine how likely it is that an actor who 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 threaten the values that network neutrality rules are intended to protect, 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, discussed in Part II.C.1, 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. The draft Open Internet Rules circulated by FCC Chairman Genachowski in early December 2010 are an example of this type of approach. They banned "unreasonable discrimination," without specifying how this term should be interpreted, as discussed in Part II.C.3.

In all case-by-case approaches, whether certain discriminatory conduct violates the nondiscrimination rule is determined in future case-by-case adjudications. As Part II.C.4 shows, this creates considerable social costs. Rules in this category provide little certainty to the market, result in high costs of regulation, and tilt the playing field against those who do not have the resources to engage in long and costly regulatory proceedings. They are also unlikely to lead to decisions that adequately protect the values network neutrality rules are intended to protect. In spite of these costs, the strategic incentives of policymakers and big stakeholders are aligned in favor of such approaches, so it is not surprising that many negotiated compromise proposals favor this type of approach.

1. *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.[177]

---

based on a broader framework, as assumed in this Article, the FCC is the right agency to enforce the rules, since it is accustomed to and tasked with applying a broader public interest standard.

177. *See, e.g.*, Becker et al., *supra* note 45, at 508; Hahn et al., *supra* note 174, at 374-79; Hazlett & Wright, *supra* note 21, at 796-806; Nuechterlein, *supra* note 176, at 43-45; J.

These proposals interpret the concerns raised by proponents of network neutral-ity regulation as concerns about anticompetitive vertical leveraging or vertical foreclosure[178] and apply the framework used to evaluate vertical leveraging and vertical foreclosure claims under U.S. antitrust laws to determine whether discriminatory conduct should be banned.[179] The term "vertical leveraging" describes a situation in which a firm that has a monopoly in one market—here, a provider of Internet access service—"abuses" or "leverages" its market power in the first market to obtain an unfair[180] advantage in a second, vertically relat-ed market—for example, in the market for a specific application.[181] The term "vertical foreclosure" applies to situations in which a monopolist in a primary market—that is, a provider of Internet access service—uses its market power in the first market to deny firms in a second, vertically related market—that is, the market for a specific application—access to that second market.[182] Over the years, the views of U.S. antitrust scholars and courts towards these practices have evolved considerably. Today, U.S. antitrust law condemns vertical lever-aging 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.[183]

Gregory Sidak & David J. Teece, *Innovation Spillovers and the "Dirt Road" Fallacy: The Intellectual Bankruptcy of Banning Optional Transactions for Enhanced Delivery over the Internet*, 6 J. COMPETITION L. & ECON. 521, 562-63 (2010); Weiser, *supra* note 174, at 74-84 (proposing an "antitrust-like approach"); Christopher S. Yoo, *Network Neutrality After* Comcast*: Toward a Case-by-Case Approach to Reasonable Network Management*, *in* NEW DIRECTIONS IN COMMUNICATIONS POLICY 55, 81-83 (Randolph J. May ed., 2009) [hereinafter Yoo, *After* Comcast]; Christopher S. Yoo, *Network Neutrality, Consumers, and Innovation*, 2008 U. CHI. LEGAL F. 179, 245-47, 257-61 [hereinafter Yoo, *Consumers & Innovation*]; Yoo, *supra* note 174, at 508-17; Maureen K. Ohlhausen, Comm'r, FTC, Remarks at MaCCI Law and Economics Conference on the Future of the Internet 11-17 (Oct. 26, 2012), http://www.ftc.gov/sites/default/files/documents/public_statements/open-internet-regulating-save-unregulated-internet/121026mannheim_0.pdf.

178. *See, e.g.*, Becker et al., *supra* note 45, at 501-02, 508; Hazlett & Wright, *supra* note 21, at 796-806; Nuechterlein, *supra* note 176, at 34; Weiser, *supra* note 174, at 71-74; Ohlhausen, *supra* note 177, at 10-11.

179. 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 dis-crimination. The text focuses on what seem to be the unifying threads in the various pro-posals.

180. HERBERT HOVENKAMP, FEDERAL ANTITRUST POLICY: THE LAW OF COMPETITION AND ITS PRACTICE § 7.9, at 348 (4th ed. 2011).

181. *See* 3 PHILLIP E. AREEDA & HERBERT HOVENKAMP, ANTITRUST LAW: AN ANALYSIS OF ANTITRUST PRINCIPLES AND THEIR APPLICATION ¶ 652a, at 130 (3d ed. 2008) [hereinafter ANTITRUST LAW] (defining vertical leveraging).

182. *See* HOVENKAMP, *supra* note 180, § 10.6b2, at 462-63 (defining foreclosure).

183. 3 AREEDA & HOVENKAMP, *supra* note 181, ¶ 652b, at 134-35; HOVENKAMP, *supra* note 180, § 7.9, at 349. Tying and exclusive dealing are evaluated according to different cri-teria, but most of the behavior that concerns network neutrality proponents does not qualify as tying or exclusive dealing.

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, U.S. 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 Phillip Areeda and Herbert Hovenkamp's antitrust treatise explains,

> Even 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.[184]

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 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 America, and Apple/iSinglePayer, discussed below in Box 9: Examples of Content-Based Discrimination), none of the content providers whose content was blocked was competing with the network provider. Similarly, a network

---

184. 3 AREEDA & HOVENKAMP, *supra* note 181, ¶ 652b, at 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 15 U.S.C. § 45(a) (2013) (section 5 of the FTC Act), either. The FTC originally adopted a different view in its order against Official Airline Guides, but this decision was reversed by *Official Airline Guides, Inc. v. FTC*, 630 F.2d 920, 927-28 (2d Cir. 1980). Since then, the FTC has signaled that it no longer deems this type of behavior actionable under section 5. *See* Federal Trade Commission Comments, Docket Nos. OST-97-2881, OST-97-3014, OST-98-4775 (June 6, 2003), *available at* http://www.ftc.gov /sites/default/files/documents/advocacy_documents/ftc-comment-department-transportation -concerning-rules-governing-airline-computer-reservation/dotcomment.pdf.

In the network neutrality context, vertical integration by network providers into applications is viewed as a prerequisite for regulatory intervention. *See, e.g.*, Martin Cave & Pietro Crocioni, *Net Neutrality in Europe*, 3 COMM. & CONVERGENCE REV. 57, 65 (2011) (explaining that consumer harm from exclusion depends on Internet service providers being vertically integrated into applications); Hahn et al., *supra* note 174, at 373, 375-76 ("[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."); Sidak & Teece, *supra* note 177, at 563 (criticizing the nondiscrimination rule proposed in the FCC's Open Internet Notice of Proposed Rulemaking as overbroad because it would apply to network providers regardless of whether they are "vertically integrated into providing competing content").

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.[185] 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 in obtaining 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 9
EXAMPLES OF CONTENT-BASED DISCRIMINATION[186]

In 2005, TELUS, Canada's second-largest Internet service provider, blocked access to a website 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 website 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's 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 President 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 were widely reported. While the latter three examples are not direct examples of Internet service providers restricting content on their networks (Verizon Wireless restricted a service on its wireless mobile network, not the wireless Internet; AT&T acted in its role as a content provider, not as an Internet service provider; and Apple acted as provider of the Apple App Store), it is easy to imagine virtually identical incidents in which an Internet service provider enacts a content policy and restricts content on its network accordingly.[187]

---

185. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 264-66.

186. This paragraph is adapted from van Schewick & Farber, *supra* note 4, at 32.

187. For a more detailed description of these examples, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 266-69. On Apple/iSinglePayer, see *iSinglePayer iPhone App Censored by Apple*, LAMBDAJIVE (Sept. 26, 2009, 9:15 PM), http://lambdajive.wordpress.com/2009/09/26/isinglepayer-iphone-app-censored-by-apple; Ryan Singel, *Developer: Apple Denied Health Care App for Political Reasons*, WIRED (Sept. 28, 2009, 6:16 PM), http://www.wired.com/epicenter/2009/09/apple-denied-health-care-app-for-political-reasons-developer-says.

Second, U.S. antitrust law only condemns vertical leveraging or vertical foreclosure as monopolization or attempted monopolization under section 2 of the Sherman Act if the monopolist is reasonably capable of monopolizing the primary market or the secondary market.[188] 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 Internet access services.[189] 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.

U.S. antitrust law generally only condemns exclusionary conduct if there is a reasonable likelihood that the behavior will harm competition, not just competitors, by worsening the structure or performance of the affected market.[190] 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.[191] 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.

---

188. 3 AREEDA & HOVENKAMP, *supra* note 181, ¶ 652a, at 130-31, ¶ 652c, at 140-42.

189. 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 multichannel 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 nondiscrimination would share this view. *See, e.g.*, Hahn et al., *supra* note 174, at 377-79 (applying an antitrust framework to the BitTorrent and Vonage cases).

190. *See* 11 HERBERT HOVENKAMP, ANTITRUST LAW, *supra* note 181, ¶ 1802b, at 74 (2011) ("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.*, Becker et al., *supra* note 45, at 501-02; Hahn et al., *supra* note 174, at 377; Farber & Katz, *supra* note 45.

191. 3 AREEDA & HOVENKAMP, *supra* note 181, ¶ 652c, at 140-42.

This requirement may be difficult to meet[192]: In many cases, the market for the application that is being discriminated against will be national in scope, while the network provider's customers only make up a part of the nation's Internet access customers.[193] For example, in the United States, the four largest broadband Internet access providers currently serve 25% (Comcast), 19% (AT&T), 14% (Time Warner), and 11% (Verizon) of the nation's broadband Internet access customers.[194] Whether a network provider's discriminatory behavior will be capable of driving the application from the market or 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 (1) the number of foreclosed Internet access customers relative to the overall number of Internet access customers, (2) the size of any economies of scale in the market for the application, and (3) the size of the cost disadvantage associated with operating at a less than efficient scale.[195] 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,[196] exclusion from access to one Internet service provider's customers may not create the type of anticompetitive harm that antitrust law is concerned about.[197] In such a case, an antitrust framework would not classify the exclusionary conduct as socially harmful.

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

---

192. *See, e.g.*, Verizon Broadband Industry Practices Comments, *supra* note 34, at 49-52; Hahn et al., *supra* note 174, at 371-72, 376; Litan & Singer, *supra* note 94, at 556-57; Sidak, *supra* note 33, at 470, 472-73; Sidak & Teece, *supra* note 177, at 566-67; Yoo, *supra* note 45, at 71-73. This question is often discussed in the context of network providers' ability to foreclose applications. *See, e.g.*, Hahn et al., *supra* note 174, at 371-72.

193. 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.

194. This data is current as of the end of the first quarter of 2014. Market share calculations are based on figures from the Leichtman Research Group. *See* Press Release, Leichtman Research Grp., Nearly 1.2 Million Add Broadband in the First Quarter of 2014 (May 20, 2014), http://www.leichtmanresearch.com/press/052014release.html.

195. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 226-32.

196. *Id.* at 231-32.

197. *See* Hahn et al., *supra* note 174, at 376; *see also* Hemphill, *supra* note 32, at 156-57; Litan & Singer, *supra* note 94, at 556; Sidak, *supra* note 33, at 470, 472-73; Hal J. Singer & J. Gregory Sidak, *Vertical Foreclosure in Video Programming Markets: Implications for Cable Operators*, 6 REV. NETWORK ECON. 372, 391-92 (2007) ("[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."); Yoo, *supra* note 45, at 71-73.

does not monopolize the market for the application in question.[198] 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.[199] 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.[200] 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.[201] The resulting decline in the amount and quality of application innovation limits the Internet's value for users and its ability to contribute to economic growth.[202] Discrimination not only deprives all Internet users of the value of future applications that would have been developed but for the threat of discrimination, but 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 (for example, 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.[203] All of these harms arise even if the behavior is unlikely to monopolize the market for the application in question.

Third, U.S. 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.[204] By contrast, network neu-

---

198. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 251-55, 264-66, 270; Frischmann & van Schewick, *supra* note 36, at 412-16.

199. *See supra* note 36 and accompanying text.

200. *See supra* Box 3; *supra* notes 59-60 and accompanying text.

201. *See infra* notes 235-36 and accompanying text.

202. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 356-61.

203. *See, e.g.*, Balkin Remarks, *supra* note 44; VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 364-65.

204. For sources stressing the importance of market power in the market for Internet services as a prerequisite for regulatory intervention, see, for example, Comments of AT&T Inc., *supra* note 91, at 66; Becker et al., *supra* note 45, at 505; Cave & Crocioni, *supra* note 184, at 65; Hahn et al., *supra* note 174, at 371; Hazlett & Wright, *supra* note 21, at 809 (arguing against FCC regulation on the ground that "market power [in the market for Internet services] is a necessary condition for such [anticompetitive] foreclosure"); Litan & Singer, *supra* note 94, at 552-54; Sidak & Teece, *supra* note 177, at 564-65; Yoo, *supra* note 174, at 511-15; and Cave et al., *supra* note 22, at 1-2. A few proponents of an antitrust framework

trality 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.[205]

Fourth, under an antitrust framework, discriminatory conduct that is justified by a legitimate business purpose would be classified as socially beneficial.[206] 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.[207] For example, most proponents of an antitrust framework seem to assume that any discriminatory conduct that is adopted to manage congestion is procompetitive and should be considered socially beneficial discrimination.[208] Price discrimination that is designed to recover fixed costs of net-

---

for network neutrality do not require proof of monopoly power or of a dominant position in the primary market. *See, e.g.*, Hahn et al., *supra* note 174, at 367.

205. *See, e.g.*, FRISCHMANN, INFRASTRUCTURE, *supra* note 106, at 330-32; VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 255-64; Christiaan Hogendorn, *Spillovers and Network Neutrality*, *in* REGULATION AND THE PERFORMANCE OF COMMUNICATION AND INFORMATION NETWORKS 191, 203-04 (Gerald R. Faulhaber et al. eds., 2012); Tim Wu, *Why Have a Telecommunications Law? Anti-Discrimination Norms in Communications*, 5 J. ON TELECOMM. & HIGH TECH. L. 15, 25, 27-28 (2006); *see also infra* Part II.D.1. In the Open Internet proceeding, this position was supported by several groups. *See, e.g.*, Center for Democracy & Technology Comments, *supra* note 36, at 9-10; Free Press Open Internet Comments, *supra* note 3, at 45-49; Open Internet Coalition Comments, *supra* note 36, at 70-76; Public Interest Comments, *supra* note 36, at 23-24.

206. *Cf.* 3 AREEDA & HOVENKAMP, *supra* note 181, ¶ 658f, at 183-92 (discussing the question of when business justification defenses will be allowed in the context of challenges under section 2 of the Sherman Act to companies' unilateral acts).

207. *See, e.g.*, Hahn et al., *supra* note 174, at 375-76, 378; Weiser, *supra* note 174, at 75-76; Christopher S. Yoo, *Network Neutrality and the Economics of Congestion*, 94 GEO. L.J. 1847, 1885-87 (2006).

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 Christopher Yoo would assume that discriminatory behavior has a procompetitive explanation unless harm to consumers is proven. *See, e.g.*, Yoo, *supra* note 45, at 66-67; Yoo, *Consumers & Innovation*, *supra* note 177, at 257-61. By contrast, Philip Weiser assumes that discriminatory conduct is anticompetitive unless a legitimate business justification is "explained" or "offered," although it is not clear what exactly would have to be proven. *See, e.g.*, Robert D. Atkinson & Philip J. Weiser, *A Third Way on Network Neutrality*, NEW ATLANTIS, Summer 2006, at 47, 57-58 (discussing discriminatory provision of Quality of Service to content providers for a fee); Philip J. Weiser, *The Next Frontier for Network Neutrality*, 60 ADMIN. L. REV. 273, 313-18 (2008) (same); Weiser, *supra* note 174, at 75-76 (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* 3 AREEDA & HOVENKAMP, *supra* note 181, ¶ 658f, at 183-85.

208. *See, e.g.*, Jerry Brito et al., Net Neutrality Regulation: The Economic Evidence at 23, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices,

work infrastructure or network innovation is often mentioned as another exam-
ple of a business justification that may legitimize discriminatory conduct.[209]
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 alter-
native that might reach the same goal with less harm to competition.[210]

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.[211] Thus, the existence of a private effi-
ciency 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 deci-
sions often diverge.[212] From the perspective of network neutrality proponents,
this divergence between the public's interests and the network providers' pri-
vate interests is a key justification for regulatory intervention. According to
them, network neutrality regulation is needed precisely because what is private-
ly 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.[213]

WC Docket No. 07-52 (Apr. 12, 2010), *available at* http://apps.fcc.gov/ecfs/document/view
?id=7020408754; *see also* Yoo, *supra* note 207, at 1907 ("An examination of the economics
of congestion provides policy justifications for precisely the type of restrictions that network
neutrality would condemn.").

209. *See, e.g.*, Howard A. Shelanski, *Network Neutrality: Regulating with More Ques-
tions than Answers*, 6 J. ON TELECOMM. & HIGH TECH. L. 23, 30 (2007); Weiser, *supra* note
207, at 315.

210. In the network neutrality context, see, for example, Hahn et al., *supra* note 174, at
375-76, 378. *See also* 3 AREEDA & HOVENKAMP, *supra* note 181, ¶ 658f, at 189-92 (describ-
ing the absence of balancing of social benefits and competitive harms under section 2 of the
Sherman Act and arguing against searching for a less restrictive alternative in some section 2
cases). *But see* United States v. Microsoft Corp., 253 F.3d 34, 59 (D.C. Cir. 2001) (requiring
the plaintiff to "demonstrate that the anticompetitive harm of the conduct outweighs the
procompetitive benefit" under section 2).

211. *See, e.g.*, FRISCHMANN, INFRASTRUCTURE, *supra* note 106, at 348-53 (discussing
price discrimination); VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at
273-78 (discussing discriminatory pricing strategies); Frischmann & van Schewick, *supra*
note 36, at 397-98, 400-07 (discussing discriminatory network management); Wu, *Network
Neutrality*, *supra* note 40, at 168-69 (discussing price discrimination); *see also* Joseph Far-
rell, *Open Access Arguments: Why Confidence Is Misplaced*, *in* NET NEUTRALITY OR NET
NEUTERING: SHOULD BROADBAND INTERNET SERVICES BE REGULATED? 195, 199-201 (Thom-
as M. Lenard & Randolph J. May eds., 2006) (discussing collateral damage from price dis-
crimination that limits application innovation).

212. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 355-75.

213. By contrast, in the context of section 2 of the Sherman Act, a legitimate business
justification only needs to be privately efficient. *See* 3 AREEDA & HOVENKAMP, *supra* note
181, ¶ 658f, at 186 ("[W]hen courts speak of the business justification defense as requiring

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 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, and users' inability to use the Internet in the way that is most valuable to them or 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.[214] 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 of 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 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-term efficiencies more important than the social benefits associated with an open, nondiscriminatory Internet, they would not have adopted network neutrality rules in the first place.

All of this does not mean that proponents of network neutrality will never allow discriminatory conduct that is motivated by considerations of private efficiency. For example, there are circumstances in 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.[215] As a re-

---

some showing of 'efficiency,' that term should be understood to refer to the costs or output of the monopolist itself (productive efficiency), not to the market as a whole (allocative efficiency)." (footnote omitted)).

214. On this trade-off, see notes 46-50 and accompanying text above.

215. To prevent such an exception from creating a loophole, any exception for reasonable network management needs to be carefully defined. *See* Marvin Ammori, *A Guide to the Network Neutrality Discussions at the FCC*, HUFFINGTON POST (Aug. 4, 2010, 3:08 PM EDT), http://www.huffingtonpost.com/marvin-ammori/a-guide-to-the-network-ne_b_670784 .html (discussing the different options for introducing loopholes into network neutrality rules); *see also* Comments of Jack Balkin et al., Preserving the Open Internet, GN Docket

sult, discriminatory conduct may be considered socially harmful by proponents of network neutrality, even if it is justified by a legitimate business justification and therefore would be allowed under an antitrust framework.[216]

In sum, a nondiscrimination rule based on an antitrust framework will not prohibit all instances of discrimination that threaten the values that network neutrality rules are designed to protect and should therefore be rejected.[217]

2. *Ban discrimination that is anticompetitive or harms users*

Other proposals would ban discrimination that is "anticompetitive" or "harms users." The proposed nondiscrimination rule may define certain behaviors as presumptively allowed or not allowed. For example, user-controlled prioritization may be presumptively legal, and 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."[218] (It included, however, an exception for reasonable network man-

No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Nov. 2, 2009), *available at* http://apps.fcc.gov/ecfs/document/view?id=7020244236 (voicing concerns regarding ambiguities in the definition of the reasonable network management exception in the Open Internet Notice of Proposed Rulemaking). For the reasonable network management exception proposed by this Article, see Box 18 and notes 491-505 and accompanying text below.

216. 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 nondiscrimination rule and exception for reasonable network management proposed by this Article. *See infra* Part II.D.2.b.

217. This conclusion is a consequence of this Article's decision to evaluate network neutrality rules based on the broad theoretical framework used by proponents of network neutrality and by the FCC in its Open Internet Order. In the context of that framework, the antitrust-based rule discussed in this Part is underinclusive. By contrast, those who approach the network neutrality debate from an antitrust framework would view the narrow scope of antitrust-based nondiscrimination rules as a "feature," not a "bug." *See* Hazlett & Wright, *supra* note 21, at 805 ("The FCC and net neutrality proponents often argue that antitrust analysis might not prohibit all use of vertical contracts is a bug, rather than a feature, of that regime. However, the fact that antitrust is not a 'slam dunk' can also be a feature." (footnote omitted)).

218. Verizon & Google, *Verizon-Google Legislative Framework Proposal*, GOOGLE 1, http://static.googleusercontent.com/media/www.google.com/en/us/googleblogs/pdfs/verizon _google_legislative_framework_proposal_081010.pdf (last visited Jan. 7, 2015) ("Non-Discrimination Requirement: In providing broadband Internet access service, a provider would be prohibited from engaging in undue discrimination against any lawful Internet con-

agement that allowed network providers "to prioritize general classes or types of Internet traffic, based on latency.")[219] The FCC-led industry negotiations in the summer of 2010 seem to have focused on a nondiscrimination rule of this type as well.[220]

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 threatens the values that network neutrality rules are designed to protect, 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.[221] 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, proponents of network neutrality 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 threaten the values that network neutrality rules are intended to protect than an interpretation based on antitrust law.[222]

---

tent, 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 nondiscrimination standard, but the presumption could be rebutted." (bolding and underlining omitted)).

219. *Id.*

220. This statement is based on conversations with various participants in the negotiations. On the FCC-led negotiations, see, for example, Cecilia Kang, *FCC Draws Fire over Talks with Internet, Telecom Giants on 'Net Neutrality,'* WASH. POST (Aug. 5, 2010), http://wapo.st/1wHWLLu [hereinafter Kang, *FCC Draws Fire*]; Cecilia Kang, *FCC Ends Talks for Deal on Net Neutrality*, WASH. POST (Aug. 6, 2010), http://wapo.st/1FEShib; Todd Shields, *Critics Decry 'Secret Deal' as AT&T, Google Huddle with FCC*, BLOOMBERG (Aug. 3, 2010), http://www.bloomberg.com/news/2010-08-02/google-verizon-at-t-seek-internet-peace-in-closed-door-talks-with-fcc.html.

221. For one example, see Farber & Katz, *supra* note 45.

222. *See supra* Part II.C.1.

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 to 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. Proponents of network neutrality 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.[223]

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

First, 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" refers only to end users or also to application and content providers.

Second, 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 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.[224]

---

223. *See, e.g.*, Barbara van Schewick, Outline of Proposal for Open Internet Rules at 2, Attachment to Notice of Ex Parte Conversations, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (May 25, 2010), *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=7020493599; Open Internet Coalition Comments, *supra* note 36, at 49 (proposing a test that requires the network management practices to result in "as little discrimination or preference as reasonably possible"); Public Interest Comments, *supra* note 36, at 40 (same); van Schewick, Official Testimony, *supra* note 68, at 4-7. The CRTC adopted a similar test to evaluate the Internet traffic management practices of Canadian Internet service providers. Review of the Internet Traffic Management Practices of Internet Service Providers, Telecom Regulatory Policy CRTC 2009-657 ¶ 43 (2009), *available at* http://www.crtc.gc.ca/eng/archive/2009/2009-657.htm (asking, among other questions, whether a discriminatory network management practice results "in discrimination or preference as little as reasonably possible").

224. *See, e.g.*, David L. Cohen, Executive Vice President, Comcast Corp., Statement at the Federal Communications Commission's Public En Banc Hearing on "Broadband Network Management Practices" at 15, Broadband Industry Practices, WC Docket No. 07-52, (Feb. 25, 2008), *available at* http://www.fcc.gov/broadband_network_management/022508 /cohen.pdf (arguing that interfering with peer-to-peer file-sharing applications "make[s] the aggregate online service better for all users and all services"); Comments of Bell Aliant Re-

Third, does it matter that there are alternative, nondiscriminatory ways of managing the network that are not similarly harmful to the users and the providers of the file-sharing software yet maintain 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 nondiscriminatory way,[225] 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.[226] Nonprofits can use peer-to-peer file sharing to distribute their video contributions to political debates.[227] Thus, peer-to-peer file-sharing applications help foster a more decentralized environment for democratic discourse and cultural production in which anybody can participate.[228] 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 concerns network neutrality proponents. However, contrary to a nondiscrimination rule based on an antitrust framework, it is at least possible that challenges to discriminatory conduct that proponents of network neutrality consider harmful will be successful.

---

gional Communications, L.P. & Bell Canada ¶¶ 86-87, Review of the Internet Traffic Management Practices of Internet Service Providers, CRTC 2008-19 (Feb. 23, 2009) (Can.) [hereinafter Bell Aliant Comments], *available at* http://www.crtc.gc.ca/public/partvii/2008 /8646/c12_200815400/1029804.zip (arguing that throttling peer-to-peer file-sharing applications between 4:30 PM and 2:00 AM "is aimed at delivering a more positive and better experience on the network for all users").

225. *See, e.g.*, *infra* Box 18; *see also infra* Part II.D.2.b.i.B.

226. Comments of the Canadian Film & Television Production Ass'n ¶¶ 53-54, 56, Review of the Internet Traffic Management Practices of Internet Service Providers, CRTC 2008-19 (Feb. 23, 2009) [hereinafter Canadian Film Comments], *available at* http://www .crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1030120.pdf (citing concrete examples); Comments of the Documentary Organization of Canada at 1-2, Review of the Internet Traffic Management Practices of Internet Service Providers, CRTC 2008-19 (Feb. 23, 2009) [hereinafter Documentary Organization Comments], *available at* http://www.crtc.gc.ca /public/partvii/2008/8646/c12_200815400/1030141.pdf (same); Comments of the Independent Film & Television Alliance ¶ 8, Review of the Internet Traffic Management Practices of Internet Service Providers, CRTC 2008-19 (Feb. 23, 2009), *available at* http://www.crtc.gc .ca/public/partvii/2008/8646/c12_200815400/1032007.doc.

227. Brief Amicus Curiae of Professors Jack M. Balkin, Jim Chen, Lawrence Lessig, Barbara van Schewick, & Timothy Wu Urging that the FCC's Order Be Affirmed at 33-35, Comcast Corp. v. FCC, 600 F.3d 642 (D.C. Cir. 2010) (No. 08-1291) [hereinafter Law Professor Comcast Amicus Brief], *available at* http://cyberlaw.stanford.edu/files/publication /files/vanschewick-2009-amicus-brief.pdf.

228. *See id.* at 31-38; Balkin, *Future of Free Expression*, *supra* note 44, at 432, 436-40.

3. *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 providers of wireline broadband Internet access without specifying how the term should be interpreted.[229] The Chairman's proposal was based on a compromise bill that had been negotiated by the Chairman of the House Committee

---

229. The draft rules were not released publicly, but they were 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." Julius Genachowski, Chairman, FCC, Remarks on Preserving Internet Freedom and Openness (Dec. 1, 2010), http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-303136A1.pdf.

While this language mirrors that of the nondiscrimination rule for common carriers in section 202 of the Communications Act (47 U.S.C. § 202(a) (2013), 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. Moreover, 47 U.S.C. § 153(51) prohibits the FCC from imposing common-carrier-type rules on entities that, like Internet service providers, have not been classified as common carriers. Verizon v. FCC, 740 F.3d 623, 649-50 (D.C. Cir. 2014). Even if interpreters were willing to look to Title II's common carrier nondiscrimination rule for guidance on how to interpret terms in a network neutrality nondiscrimination 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. Harold Feld, *Genachowski Hits the Legal Reset Button—"Title II Lite*," WETMACHINE (May 7, 2010), http://tales-of-the-sausage-factory.wetmachine.com/content/genachowski-hits-the-legal-reset-button-title-ii-lite; Harold Feld, *Sorry AT&T, Title II Would Not Require the FCC to Allow Paid Prioritization*, WETMACHINE (Oct. 8, 2010), http://www.wetmachine.com/content/sorry-att-title-ii-would-not-require-the-fcc-to-allow-paid-prioritization; Hank Hultquist, *Harold Feld is Right (About Some Things)*, AT&T PUB. POL'Y BLOG (Oct. 13, 2010, 5:32 PM), http://attpublicpolicy.com/government-policy/harold-feld-is-right-about-some-things; Bob Quinn, *Who Keeps Pulling the Net Neutrality Football?*, AT&T PUB. POL'Y BLOG (Sept. 9, 2010, 10:00 AM), http://attpublicpolicy.com/government-policy/who-keeps-pulling-the-net-neutrality-football; *see also* Daniel A. Lyons, *Net Neutrality and Nondiscrimination Norms in Telecommunications*, 54 ARIZ. L. REV. 1029, 1058-64 (2012).

Section 202(a) reads in full:

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.

47 U.S.C. § 202(a).

on Energy and Commerce, Representative Henry A. Waxman, and the Chairman of the House Subcommittee on Communications, Technology and the Internet, Representative Rick Boucher, with the large phone and cable network providers, Internet companies, consumer groups, and open Internet groups in the fall of 2010.[230] 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."[231]

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 decisionmakers with maximum flexibility. Contrary to nondiscrimination rules based on an antitrust framework, this type of proposal does not immediately rule out cases that concern network neutrality proponents and makes it at least possible, but not certain, that a complaint targeting behavior that network neutrality proponents deem socially harmful will be successful.

4. *Problems with case-by-case adjudication*

All of the proposals in this Subpart leave the substantive decision over the legality of specific discriminatory behavior to future case-by-case adjudications. The most general proposals ban "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 and their application to the specific behavior under consideration to future decisionmakers. While proposals that prohibit discrimination that "causes 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 inter-

---

230. *See* Press Release, Henry A. Waxman, Waxman Statement on Net Neutrality Proposal (Sept. 29, 2010), http://waxman.house.gov/chairman-waxmans-statement-net -neutrality-proposal; *see also* Sara Jerome, *Waxman May Move Without Full Consensus on Net-Neutrality Bill*, HILL (Sept. 23, 2010, 5:32 PM ET), http://thehill.com/policy/technology /120571-waxman-may-move-without-full-consensus-on-net-neutrality-bill. Not all participants in the negotiations supported the final proposal. John Eggerton, *Free Press Pushed Hard Against Waxman's Net Neutrality Bill*, MULTICHANNEL NEWS (Sept. 30, 2010), http:// www.multichannelnews.com/article/457836-Free_Press_Pushed_Hard_Against_Waxman _s_Net_Neutrality_Bill.php; Sara Jerome, *Sources: OIC Not Supporting Waxman Net-Neutrality Bill*, HILL (Sept. 28, 2010, 10:39 PM EDT), http://thehill.com/policy/technology /121493-sources-oic-will-not-support-waxman-net-neutrality-bill [hereinafter Jerome, *OIC*]. It was dropped when the Republican members of the House Committee on Energy and Commerce refused to support the bill. Press Release, Henry A. Waxman, *supra*; *see also* Kim Hart, *Net Neutrality Bill Stillborn*, POLITICO (Sept. 29, 2010, 6:47 PM EST), http:// www.politico.com/news/stories/0910/42919.html. Chairman Waxman later filed the draft bill in the Open Internet docket. Open Internet Act of 2010 Draft Bill, Attachment to Comments of Henry A. Waxman, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 1, 2010) [hereinafter Draft Bill], *available at* http://apps.fcc.gov/ecfs/document/view?id=7020922223. On the Waxman proposal's influence on the Chairman's proposal, see notes 546, 558-60 and accompanying text below.

231. Draft Bill, *supra* note 230, § 2, at 2 (proposing to add § 12(a)(1)(B)).

pretations leading to radically different outcomes. Other nondiscrimination 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 decisionmakers interpret and apply those factors, but also on how they weigh the different factors against each other. The nondiscrimination rule proposed by the FCC in its May 2014 Notice of Proposed Rulemaking is an example of such a rule.[232]

These kinds of case-by-case approaches create considerable social costs.[233]

a. *Lack of certainty and predictability*

First, case-by-case approaches fail to provide much-needed certainty for industry participants.

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.[234] 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.

More generally, some research and anecdotal evidence suggest that in the broadband context, certainty regarding the regulatory framework and its stabil-

---

232. The nondiscrimination 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. 29 FCC Rcd. 5561, 5599-610, 5618-21 (proposed May 15, 2014) (to be codified at 47 C.F.R. pt. 8).

233. 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. the rule of reason in antitrust law.

234. Comments of the National Cable & Telecommunications Ass'n, *supra* note 91, at 15; Verizon Broadband Industry Practices Comments, *supra* note 34, at 44-45.

ity over time may be more important for network investment than the substance of the regulatory decision.[235]

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.[236] 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 in which 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 nondiscrimination 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 nondiscrimination rule does not sufficiently protect users and application developers against actual

---

235. *See, e.g.*, Verizon Broadband Industry Practices Comments, *supra* note 34, at 44-45; CHRISTIAN WERNICK, STRATEGIC INVESTMENT DECISIONS IN REGULATED MARKETS: THE RELATIONSHIP BETWEEN INFRASTRUCTURE INVESTMENTS AND REGULATION IN EUROPEAN BROADBAND 129-30, 158-85 (2007).

236. *See, e.g.*, *Hearing on the Future of the Internet*, *supra* note 33, at 55-56 (prepared statement of Lawrence Lessig, Professor of Law, Stanford Law School); VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 270-73.

As I know from many conversations with entrepreneurs and investors, the threat of discrimination reduces entrepreneurs' ability to secure funding today. For two publicly documented examples, see Barbara van Schewick, Assistant Professor of Law, Stanford Law Sch., Oral Testimony at the Federal Communications Commission's Second Public En Banc Hearing on Broadband Management Practices at 2, Broadband Industry Practices, WC Docket No. 07-52 (Apr. 17, 2008) [hereinafter van Schewick, Oral Testimony], *available at* www.fcc.gov/broadband_network_management/041708/vanschewick-oral.pdf, and the letter submitted by the founders of the online video company Zediva to the FCC, Ex Parte Letter of Zediva at 1-2, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 10, 2010) [hereinafter Zediva Ex Parte Letter], *available at* http://apps.fcc.gov/ecfs/document/view?id=7020923207. *See also Hearing on the Future of the Internet*, *supra* note 33, at 55-56 (prepared statement of Lawrence Lessig, Professor of Law, Stanford Law School).

discrimination and fails to remove the threat of discrimination as a factor that affects application developers' and innovators' decisions about innovation.[237]

While individual adjudications may reduce the amount of uncertainty over time, it is unclear whether and how fast useful precedents will emerge.

Over time, individual adjudications may clarify the interpretation of the standard and its application to specific behavior, reducing uncertainty.[238] 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.[239] Thus, it is unclear whether and how quickly useful precedents will emerge. In the meantime, the costs associated with the uncertainty persist.[240] Moreover, as set out in more detail below, the substantive principles emerging from case-by-case adjudications are

---

237. 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 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.

238. See, e.g., Yoo, *After* Comcast, *supra* note 177, at 82.

239. See, e.g., Louis Kaplow, *Rules Versus Standards: An Economic Analysis*, 42 Duke L.J. 557, 612-13 (1992) (discussing courts' tendency to issue narrow decisions that minimize or postpone the creation of precedents); Cass R. Sunstein, *Incompletely Theorized Agreements*, 108 Harv. L. Rev. 1733, 1735-38, 1739-51 (1995) (discussing the dynamics that give rise to this phenomenon).

240. See, e.g., Yoo, *After* Comcast, *supra* note 177, at 57, 82 (noting the uncertainty surrounding the FCC's interpretation of "reasonable network management" in the aftermath of the FCC's order against Comcast until enough precedents exist). *See generally* Louis Kaplow, *General Characteristics of Rules*, *in* 5 Encyclopedia of Law and Economics: The Economics of Crime and Litigation 502, 512 (Boudewijn Bouckaert & Gerrit De Geest eds., 2000) (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); Kaplow, *supra* note 239, at 612-14.

less likely to adequately protect the values and actors that network neutrality rules are designed to protect.

b. *High costs of regulation*

Second, case-by-case approaches create high costs of regulation.[241] 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.[242] Their applicability may also be limited by the facts of the case.[243] 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,[244] 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 (IETF) standards for the operation of the TCP.[245] 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.[246] Third, Comcast had not disclosed the use of the practice to its Internet access customers.[247] 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, of the criteria used in the Comcast case.[248]

---

241. *See, e.g.*, 1 RICHARD J. PIERCE, JR., ADMINISTRATIVE LAW TREATISE §§ 6.8-.9, at 497-508 (5th ed. 2010); Kaplow, *supra* note 240, at 510.

242. *See* 1 PIERCE, *supra* note 241, § 6.8, at 498-99; M. Elizabeth Magill, *Agency Choice of Policymaking Form*, 71 U. CHI. L. REV. 1381, 1394, 1396 (2004).

243. *See* 1 PIERCE, *supra* note 241, § 6.8, at 499-500; Magill, *supra* note 242, at 1396.

244. Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13,028, 13,059-61 (2008) (memorandum opinion and order), *vacated*, Comcast Corp. v. FCC, 600 F.3d 642 (D.C. Cir. 2010).

245. *Id.* at 13,054-55.

246. *Id.* at 13,055-58.

247. *Id.* at 13,058-59.

248. The Comcast Order was vacated by the U.S. Circuit Court of Appeals for the District of Columbia in April 2010. *Comcast*, 600 F.3d at 661.

c. *Limited ability to protect values and actors that network neutrality
   rules are designed to protect*

Finally, in the context of network neutrality, case-by-case approaches are
less likely than rule-based approaches to adequately protect the values and ac-
tors 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,
start-ups, nonprofits, independent artists, and members of underserved commu-
nities—who do not have the resources necessary to engage in extended fights
over the legality of specific instances of discrimination in the future.[249] Net-
work providers and large application providers can conduct fact-intensive in-
vestigations, pay lawyers, economists, and other experts to engage in the fight
over the correct interpretation and 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 applica-
tion 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.[250]

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 il-
luminate all relevant aspects of the underlying problem may distort the
decisionmaker's view of the underlying policy issues in a way that a more gen-
eral analysis of the issues in the context of a rulemaking proceeding may
not.[251] 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 do not know that
peer-to-peer file-sharing applications have many legal and socially valuable us-
es. For example, at the time of Comcast's blocking of BitTorrent, established
content providers such as the BBC, Showtime, the History Channel, MTV

---

249. *See, e.g.*, Steven P. Croley, *Theories of Regulation: Incorporating the Administra-
tive Process*, 98 COLUM. L. REV. 1, 120-24, 128 (1998); Cass R. Sunstein, *Problems with
Rules*, 83 CALIF. L. REV. 953, 977 (1995) ("[C]ase-by-case judgments systematically favor
the well-to-do."). In the network neutrality context, see van Schewick, August 2010 Attach-
ment, *supra* note 52, at 5; and Brad Burnham, Op-Ed., *A Threat to Startups*, N.Y. TIMES
ROOM FOR DEBATE (Aug. 10, 2010, 12:08 PM), http://www.nytimes.com/roomfordebate
/2010/8/9/who-gets-priority-on-the-web/a-threat-to-startups.

250. On the importance of low-cost innovators, see VAN SCHEWICK, ARCHITECTURE AND
INNOVATION, *supra* note 2, at 204-13, 300-08, 334-45. On new entrants and start-ups, see *id.*
at 319-34. For a short version of the argument, see van Schewick, Open Internet Opening
Statement, *supra* note 29, at 3-5.

251. *See* 1 PIERCE, *supra* note 241, at 496-97; Magill, *supra* note 242, at 1396; *see also*
Kaplow, *supra* note 239, at 609 (discussing this phenomenon with respect to the decisions of
courts compared to legislative decisions).

Networks, 20th Century Fox, and Paramount were distributing their video content online through services that utilized the BitTorrent protocol.[252] Developers of open source applications such as the Linux operating system or OpenOffice and 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.[253]

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.[254] 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.[255] Still, based on the inaccurate perception that applications like BitTorrent are primarily used for illegal file sharing, regulators and 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 stewards 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, BT restricted the bandwidth available to the BBC iPlayer and other streaming video applications to 896 kilobits per second in a particular version of BT's broadband service.[256] 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 thirty-four percent of downstream traffic on fixed broadband networks during

---

252. Formal Complaint at 17-18, Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications, 23 FCC Rcd. 13,028 (2008) (No. EB-08-IH-1518).

253. *Id.* at 18-19; LINUXTRACKER, http://linuxtracker.org (last visited Jan. 7, 2015); *Legacy OpenOffice.org P2P Downloads*, OPENOFFICE, http://www.openoffice.org /distribution/p2p (last visited Jan. 7, 2015).

254. Canadian Film Comments, *supra* note 226, at ¶¶ 53-56, 58; Documentary Organization Comments, *supra* note 226, at 1-3, 5.

255. *See* Law Professor Comcast Amicus Brief, *supra* note 227, at 31-38.

256. Cellan-Jones, *supra* note 101.

peak times.[257] As a result, in a generalized rulemaking that also considers limits on applications other than peer-to-peer file-sharing applications, the sympathy of decisionmakers and observers will be more evenly distributed among restricted and unrestricted uses of the network. At the same time, streaming applications, which compete with network providers' traditional video offerings, bring the potential gap between network providers' and users' interests into sharp relief,[258] 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 than a general rulemaking to result in a decision that grants network providers broad discretion in managing congestion. 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.[259] In contrast 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.[260] 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

---

257. SANDVINE, GLOBAL INTERNET PHENOMENA REPORT 1H 2014, at 6 tbl.2 (2014), *available at* https://www.sandvine.com/downloads/general/global-internet-phenomena/2014/1h-2014-global-internet-phenomena-report.pdf.

258. *See, e.g.*, Competitive Impact Statement at 11, 14-20, United States v. Comcast Corp., 808 F. Supp. 2d 145 (D.D.C. 2011) (No. 1:11-cv-00106) [hereinafter Competitive Impact Statement], *available at* http://www.justice.gov/atr/cases/f266100/266158.pdf (discussing the gap and citing "[m]any internal documents" showing that Comcast views online video distributors as a competitive threat to its traditional cable video distribution offerings); Canadian Film Comments, *supra* note 226, ¶¶ 60-67.

259. *See generally* 1 PIERCE, *supra* note 241, § 6.8, at 496-97, 501; Magill, *supra* note 242, at 1396, 1446.

260. *See generally* 1 PIERCE, *supra* note 241, § 6.8, at 496-97, 501; Croley, *supra* note 249, at 116-17, 120-24, 128, 148; Magill, *supra* note 242, at 1391, 1396.

proceeding.[261] Adversarial proceedings increase these challenges.[262] For example, entrepreneurs are often reluctant to speak out on network neutrality because they fear retaliation by network providers.[263] They may be even more reluctant to do so in the context of an adjudication that 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 decisionmakers, they may also be subject to biases or intuitive reactions resulting from an adjudication's specific fact patterns.[264] 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 decisionmaker of input from important stakeholders.

In addition, 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.[265] 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 create 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 create social costs. Thus, there is a trade-off that regulators need to resolve. An ex ante rule that specifies what be-

---

261. MANCUR OLSON, THE LOGIC OF COLLECTIVE ACTION: PUBLIC GOODS AND THE THEORY OF GROUPS 2, 9-16, 35 (1965); Nicholas Bagley & Richard L. Revesz, *Centralized Oversight of the Regulatory State*, 106 COLUM. L. REV. 1260, 1287-90 (2006); Rachel E. Barkow, *Insulating Agencies: Avoiding Capture Through Institutional Design*, 89 TEX. L. REV. 15, 21-24 (2010); Croley, *supra* note 249, at 126-42.

262. Croley, *supra* note 249, at 120-24, 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 decisionmaking, as well as citing empirical studies to support his analysis).

263. Barbara van Schewick, *Start-Up Video Company Asks FCC to Improve Open Internet Proposal*, INTERNET ARCHITECTURE & INNOVATION (Dec. 12, 2010, 9:47 PM), http://netarchitecture.org/2010/12/start-up-video-company-files-concerns-about-fcc-open-internet-proposal.

264. *See supra* notes 251-58 and accompanying text; *infra* notes 268-78, 303-04 and accompanying text.

265. For a detailed discussion of this trade-off, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 355-71. For a short overview, see the discussion in Part I above.

havior 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.[266]

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.[267]

Moreover, an adjudicator is likely to underestimate other negative consequences of allowing a deviation from the general nondiscrimination rule in the particular case under 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.[268] 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[269]: 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.[270] The increase in encryption has motivated some operators to slow down all encrypted traffic, which in turn

---

266. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 78 (making the same argument in the context of deviations from the broad version of the end-to-end arguments); *see also* Lawrence B. Solum & Minn Chung, *The Layers Principle: Internet Architecture and the Law*, 79 NOTRE DAME L. REV. 815, 854-65 (2004) (making a similar argument in a slightly different context).

267. For a more detailed description of the problem with citations to relevant literature, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 77-78, 374-75.

268. *See id.* at 78; Solum & Chung, *supra* note 266, at 859-60.

269. The resulting arms race is described in detail by William H. Lehr et al., *Scenarios for the Network Neutrality Arms Race*, 1 INT'L J. COMM. 607, 607-08, 614-40 (2007).

270. *See* Blumenthal & Clark, *supra* note 70, at 77-78, 95; Lehr et al., *supra* note 269, at 627-28; Karl Bode, *Which ISPs Throttle BitTorrent Traffic?*, BROADBAND DSLREPORTS.COM (Dec. 22, 2006), http://www.dslreports.com/shownews/80468.

hurts legitimate traffic that is encrypted for security reasons.[271] Widespread use of encryption also complicates network analysis, planning, and security.[272]

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 nonstandard 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.[273] 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.[274] For example, while application developers may be able to adapt their application to one network provider's idiosyncratic discriminatory network management practice, the costs of adapting their application to the network management practices of more than a few providers will quickly become prohibitive.[275] 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 and its interactions with other current or future deviations from nondiscriminatory network management. By

---

271. 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's customers and were checking their e-mail from off campus; off-campus access to e-mail is encrypted for security reasons. *See* Michael Geist, Op-Ed., *ISP Must Come Clean on 'Traffic Shaping,'* THESTAR.COM (Apr. 16, 2007), http://www.thestar.com/business/2007/04/16/isp_must_come_clean_on_traffic_shaping .html. 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. *See* Ken Engelhart, Letter to the Editor, *Cyber Traffic Moving Smoothly*, THESTAR.COM (Apr. 20, 2007, 4:30 AM), https://web.archive.org/web /20071118123710/http://www.thestar.com/article/205204 (accessed via the Internet Archive index).

272. *See* Blumenthal & Clark, *supra* note 70, at 77-78, 86; Lehr et al., *supra* note 269, at 635.

273. *See, e.g.*, PETER ECKERSLEY ET AL., ELEC. FRONTIER FOUND., PACKET FORGERY BY ISPS: A REPORT ON THE COMCAST AFFAIR 3-6, 8-9 (2007), *available at* http://www.eff.org /files/eff_comcast_report.pdf. 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.

274. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 78, 374-75; Solum & Chung, *supra* note 266, at 855-56, 859-60.

275. Letter from Lawrence Lessig, Professor of Law, Stanford Law Sch., to Marlene H. Dortch, Sec'y, FCC 2 (Aug. 20, 2008), *available at* http://cyberlaw.stanford.edu/files/blogs /LessigFccComcastLetter.pdf.

contrast, decisionmakers in a general rulemaking can take a broader view that takes account of cumulative effects and generalizes from past experiences.[276]

Finally, research in behavioral economics suggests that individuals tend to systematically undervalue future benefits, discounting them more than rational discounting would suggest.[277] Uncertainty about future benefits aggravates this bias.[278] 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.

So far, the discussion of the social costs of case-by-case proposals in this Subpart has focused on the costs associated with general or ambiguous nondiscrimination standards. Although case-by-case approaches based on an antitrust framework provide considerably more guidance on how to evaluate discriminatory behavior, the outcome of specific cases under an antitrust framework still depends on the exact interpretation of the framework and on its application to the facts. In addition, the results of cases under an antitrust framework turn on facts (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) that are highly specific to individual cases and that are often difficult and costly to prove.[279] As a result, an antitrust framework is afflicted with the same social costs as case-by-case proposals based on more general or ambiguous standards. In particular, the uncertainty about the legality of specific discriminatory conduct is not resolved until after the discrimination has occurred. In addition, 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

---

276. *See* Solum & Chung, *supra* note 266, at 855-56, 859-60.

277. This bias is known as hyperbolic discounting. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 78, 374 (discussing how hyperbolic discounting affects decisions to deviate from the broad version of the end-to-end arguments). *See generally* Shane Frederick et al., *Time Discounting and Time Preference: A Critical Review*, 40 J. ECON. LITERATURE 351, 366-67 (2002); George Loewenstein & Drazen Prelec, *Anomalies in Intertemporal Choice: Evidence and an Interpretation*, 107 Q.J. ECON. 573 (1992); Daniel Read, *Intertemporal Choice*, *in* BLACKWELL HANDBOOK OF JUDGMENT AND DECISION MAKING 424 (Derek J. Koehler & Nigel Harvey eds., 2004).

278. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 78, 374. *See generally* Read, *supra* note 277, at 431.

279. *See* 2B PHILLIP E. AREEDA ET AL., ANTITRUST LAW, *supra* note 181, ¶ 408b, at 44-48 (2007) (arguing that precise measurement of economies of scale is very difficult, if not impossible); *id.* ¶ 408d, at 54-58 (arguing that "antitrust rules requiring judicial measurement of scale economies should be avoided" because "[c]ourts are simply not up to the task of deciding *ex jure . . .* that a practice should be condemned because it denies a firm scale economies except in the clearest of cases").

necessarily remove the uncertainty. Finally, like the general or ambiguous non-discrimination rules discussed above, a nondiscrimination 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.

d.   *Strategic incentives of policymakers and big stakeholders*

In spite of these considerable social costs associated with general or ambiguous nondiscrimination standards, 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.[280] 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 United States, AT&T, Verizon, Comcast, and Time Warner) or big providers of Internet applications, content, or services such as Google and Amazon. Under these circumstances, adopting a very general or ambiguous nondiscrimination rule today constitutes an attractive compromise, since the controversial question is not decided one way or the other.[281]

The legislator or regulator can reap any immediate benefits associated with adopting network neutrality rules[282] while avoiding the immediate political

---

280. *See, e.g.*, William N. Eskridge, Jr., *Politics Without Romance: Implications of Public Choice Theory for Statutory Interpretation*, 74 Va. L. Rev. 275, 288-89 (1988); Kaplow, *supra* note 239, at 609 nn.141 & 143; Sunstein, *supra* note 249, at 973, 1004-05.

281. *See generally* Eskridge, *supra* note 280, at 288-89; Sunstein, *supra* note 249, at 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 a judge.").

282. 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. *Barack Obama: Connecting and Empowering All Americans Through Technology and Innovation*, Obama for Am., http://my.barackobama.com/page/-/HQpress/111307 %20Innovation%20fact%20sheet.pdf (last visited Jan. 7, 2015) [hereinafter *Barack Obama*]. Likewise, the FCC's adoption of network neutrality rules satisfied the Democratic leadership in Congress, which also had supported network neutrality regulation. Marvin Ammori, *All the Political Cover the FCC Could Ever Want*, Huffington Post (Sept. 30, 2010, 11:30 AM), http://www.huffingtonpost.com/marvin-ammori/all-the-political-cover-t_b_745243 .html; Harold Feld, *Genachowski's Fast Fading Star—And How He Can Still Salvage His Term as Chairman*, Wetmachine (Aug. 3, 2010), http://tales-of-the-sausage-factory .wetmachine.com/content/genachowskis-fast-fading-star-and-how-he-can-still-salvage-his-term-as-chairman. The adoption of network neutrality rules in the United States was also supported by large Internet companies (e.g., Amazon, eBay, Facebook, and Google), entrepreneurs, and investors, Open Internet Coalition Comments, *supra* note 36, at 76; Letter

costs of taking on powerful interests on one side of the debate.[283] While an ambiguous or general nondiscrimination rule that is applied case by case is more difficult and costly to apply and enforce in the future, these costs will not be borne by the entity adopting the rule.[284] If the nondiscrimination rule is adopted through legislation, it will most likely be enforced by a regulatory agency (e.g., in the United States, by the FCC). Even if the nondiscrimination rule is adopted by a regulatory agency such as the FCC through administrative rulemaking, it may be enforced by future members of the agency (e.g., in the case of the FCC, by future Commissioners) 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 nondiscrimination rule, either.[285]

Big stakeholders support this type of nondiscrimination rule because each side can claim a win (or at least a nonloss) 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 nondiscrimination rule described above emerged from industry negotiations at the FCC[286] and in Congress,[287] or, as in

---

from Immad Akhund et al. to Julius Genachowski, Chairman, FCC (Oct. 21, 2009), Open Internet Coalition Comments, *supra* note 36, app. B at 96-97; public interest groups, Letter from Helen Soule et al. to Julius Genachowski, Chairman, FCC (Oct. 21, 2009), *available at* https://www.publicknowledge.org/pdf/pubint-nn-letter-20091021.pdf; and a large grassroots movement, Timothy Karr, *Netroots to Obama FCC: Inaction Is Not an Option*, Huffington Post (July 4, 2010, 5:12 AM EDT), http://www.huffingtonpost.com/timothy-karr/netroots -to-obama-fcc-ina_b_562564.html.

283. *See, e.g.*, Kaplow, *supra* note 239, at 609 nn.141 & 143 (discussing the cost-benefit calculus described in the text in the context of a legislature's choice between rules and standards); Sunstein, *supra* note 249, at 973, 1004, 1013 (discussing the same issue).

284. *See, e.g.*, Sunstein, *supra* note 249, at 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.").

285. *See, e.g.*, *id.* at 973.

286. 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. *See* Kang, *FCC Draws Fire*, *supra* note 220.

287. The FCC Chairman's draft Open Internet Rules were based on a proposal for a network neutrality bill that had been negotiated by Representative Waxman, the Chairman of the House Committee on Energy and Commerce, and Representative 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 supra* note 230. On the influence of the draft bill on the Chairman's draft rules, see notes 546, 558-60 and accompanying text below.

the Verizon-Google legislative framework proposal, from direct negotiations between two big stakeholders on opposite sides of the debate.[288]

D. *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 the standards-based approaches, however, these proposals 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.

Out of the three proposals in this category, only one—ban application-specific discrimination, but allow application-agnostic discrimination (Part II.D.2.b)—accurately distinguishes socially beneficial from socially harmful discrimination. This is the rule that policymakers should adopt. By contrast, the other two proposals—ban discrimination that is not disclosed (Part II.D.1) and ban discrimination that does not treat like traffic alike (Part II.D.2.a)—do not adequately protect the values that network neutrality rules are intended to protect and should be rejected.

1. *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 the formal criterion of whether the network provider disclosed the differential treatment. Alternatively, a network neutrality regime might allow blocking or discrimination but require Internet service providers to disclose any blocking or discrimination that occurs.[289] In January 2014, the Court of Appeals for the D.C. Circuit struck down the Open Internet Order's rules against

---

288. 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. Press Release, Leichtman Research Grp., Over 635,000 Add Broadband in the Third Quarter of 2011 (Nov. 18, 2011), http://www.leichtmanresearch.com/press/111811release.html; *see also Market Share of Wireless Subscriptions Held by Carriers in the U.S. from 1st Quarter 2011 to 2nd Quarter 2014*, STATISTA, http://www.statista.com/statistics/199359/market-share-of-wireless-carriers -in-the-us-by-subscriptions (last visited Jan. 7, 2015) [hereinafter *Wireless Market Share 2011-2014*].

289. *See, e.g.*, Gerald R. Faulhaber, *A National Broadband Plan for Our Future: A Customer-Centric Framework*, 3 INT'L J. COMM. 742, 762-66 (2009); Faulhaber & Farber, *supra* note 45, at 315-16, 327-33, 336; John W. Mayo et al., Op-Ed., *How to Regulate the Internet Tap*, N.Y. TIMES (Apr. 20, 2010), http://www.nytimes.com/2010/04/21/opinion/21mayo .html; Cave et al., *supra* note 22, at 2-3; AT&T et al., *supra* note 22, at 1, 3.

blocking and discrimination but upheld the disclosure rule.[290] Thus, until the FCC adopts new network neutrality rules, the current network neutrality regime in the United States constitutes an example of this approach.

In 2009, the European Union adopted this approach following the review of its regulatory framework for telecommunications services.[291] The European Universal Service Directive neither requires network providers to impose restrictions on users' use of applications nor prevents them from doing so.[292] 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, and about any traffic management measures and their impact on service quality. This information must be disclosed in the terms of the contract and when practices change.

---

290. Verizon v. FCC, 740 F.3d 623, 659 (D.C. Cir. 2014).

291. *See* Directive 2002/22/EC of the European Parliament and of the Council of 7 March 2002 (Universal Service Directive), arts. 20-21, 2002 O.J. (L 108) 51, 64-65, as amended by Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009, 2009 O.J. (L 337) 11.

292. *See* Directive 2009/136/EC, *supra* note 291, pmbl. ¶ 29, at 14 ("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.

Id. At the same time, the Universal Service Directive, like the Framework Directive, stresses 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).

Id. pmbl. ¶ 28, at 14.

The Framework Directive explicitly requires national regulatory authorities to promote this goal:

> 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*: . . . ensuring that there is no distortion or restriction of competition in the electronic communications sector, including the transmission of content . . . . 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 . . . .

Directive 2002/21/EC of the European Parliament and of the Council of 7 March 2002 (Framework Directive), art. 8, 2002 O.J. (L 108) 33, 41-42, as amended by Directive 2009/140/EC of the European Parliament and of the Council of 25 November 2009, art. 1(8), 2009 O.J. (L 337) 37, 50-51.

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 that does not discriminate against the affected application. The threat of switching, proponents of this approach assume, will discipline providers.[293]

In line with this reasoning, 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 United States.[294] Similar arguments are often made for mobile Internet access, where users often have a choice between three or more competitors.[295]

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 some applications and content on its network without losing too many Internet service customers.[296] 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 compe-

---

293. *See, e.g.*, *Commission Staff Working Document: Impact Assessment*, at 92, SEC (2007) 1472 (Nov. 13, 2007) [hereinafter *Impact Assessment*], *available at* http://www .europarl.europa.eu/registre/docs_autres_institutions/commission_europeenne/sec/2007/1472 /COM_SEC%282007%291472_EN.pdf; European Commission, *The Open Internet and Net Neutrality in Europe*, at 4, COM (2011) 222 final (Apr. 19, 2011), *available at* http://eur-lex .europa.eu/LexUriServ/LexUriServ.do?uri=COM:2011:0222:FIN:EN:PDF.

294. *See Impact Assessment*, *supra* note 293, at 91-92; KENNETH R. CARTER ET AL., NETWORK NEUTRALITY: IMPLICATIONS FOR EUROPE, at v-vi, 56-57 (WIK Diskussionsbeitrag, No. 314, 2008); William J. Baumol et al., *Economists' Statement on Network Neutrality Policy* 2 (AEI-Brookings Joint Ctr. for Regulatory Studies, Related Publication 07-08, 2007), *available at* http://ssrn.com/abstract=976889.

295. *See generally* Gerald R. Faulhaber & David J. Farber, *Innovation in the Wireless Ecosystem: A Customer-Centric Framework*, 4 INT'L J. COMM. 73 (2010).

296. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 264; van Schewick, *supra* note 40, at 374-77.

tition in the market for Internet access services, disclosure cannot replace substantive regulation as a tool to discipline providers.[297]

a.   *Problems with disclosure-only network neutrality regimes*

Disclosure can only discipline providers if there is effective competition.[298] 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 cost. Even if there is competition in the market for Internet access services, these conditions will often not be met.

i.   *Consumers' incomplete knowledge, cognitive limitations, and cognitive biases*

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.[299] 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

---

297. 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, ARCHITECTURE AND INNOVATION, *supra* note 2, at 255-59. On the impact of competition on Internet service providers' incentives to engage in discriminatory traffic management, see generally Alissa Cooper, How Competition Drives Discrimination: An Analysis of Broadband Traffic Management in the UK (Aug. 2013) (unpublished manuscript) [hereinafter Cooper, Analysis], *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2241562 (basing conclusions on a case study of broadband traffic management in the United Kingdom); and Alissa Cooper, How Regulation and Competition Influence Discrimination in Broadband Traffic Management: A Comparative Study of Net Neutrality in the United States and the United Kingdom ch. 5, at 105-29, ch. 7, at 171-211 (Sept. 2013) (unpublished Ph.D. thesis, University of Oxford) [hereinafter Cooper, Thesis], *available at* http://www.alissacooper.com/files/Thesis.pdf (same).

298. The following discussion focuses on the merits of a nondiscrimination 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. On the benefits of disclosure rules as a complement to substantive network neutrality rules, see Box 11 below.

299. 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.*, Oren Bar-Gill, *Bundling and Consumer Misperception*, 73 U. CHI. L. REV. 33 (2006); Richard Craswell, *Tying Requirements in Competitive Markets: The Consumer Protection Issues*, 62 B.U. L. REV. 661 (1982).

will not necessarily be the first explanation that comes to mind.[300] Even if users consider that possibility, many will lack the expertise to investigate the cause of the bad performance.[301] 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.[302] Consumers often do not read disclosures, and in many cases, those who read them do not understand them.[303] For those who read and understand the disclosure, knowing which 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

---

300. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 260-61; van Schewick, *supra* note 40, at 376-77.

301. 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* ECKERSLEY ET AL., *supra* note 273, at 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 *Test Your ISP*, ELEC. FRONTIER FOUND., https://www.eff.org/testyourisp (last visited Jan. 7, 2015). To help foster the creation of tools that consumers can use to monitor their network connections, the FCC 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." *FCC Open Internet Apps Challenge*, CHALLENGEPOST, http://openinternetapps.challenge.gov (last visited Jan. 7, 2015). The winners were announced in August 2011. *Chairman Announces Challenge.gov Competition Winners*, CHALLENGEPOST (Aug. 5, 2011), http://openinternetapps.challengepost.com/updates/57-chairman-announces -challenge-gov-competition-winners.

302. *See* Omri Ben-Shahar & Carl E. Schneider, *The Failure of Mandated Disclosure*, 159 U. PA. L. REV. 647, 665-79, 704-29 (2011) (reviewing the experience with disclosure in a variety of contexts); Fred H. Cate, *The Failure of Fair Information Practice Principles*, *in* CONSUMER PROTECTION IN THE AGE OF THE 'INFORMATION ECONOMY' 341, 341-42 (Jane K. Winn ed., 2006) (privacy); Matthew A. Edwards, *Empirical and Behavioral Critiques of Mandatory Disclosure: Socio-Economics and the Quest for Truth in Lending*, 14 CORNELL J.L. & PUB. POL'Y 199, 204 (2005) (truth in lending); Howard Latin, *"Good" Warnings, Bad Products, and Cognitive Limitations*, 41 UCLA L. REV. 1193, 1206-22 (1994) (product warnings); Susanna Kim Ripken, *The Dangers and Drawbacks of the Disclosure Antidote: Toward a More Substantive Approach to Securities Regulation*, 58 BAYLOR L. REV. 139 (2006) (securities regulation).

303. *See, e.g.*, Ben-Shahar & Schneider, *supra* note 302, at 666, 668-69, 671-79, 709-18 (citing studies on disclosures in a variety of contexts); M. Ryan Calo, *Against Notice Skepticism in Privacy (and Elsewhere)*, 87 NOTRE DAME L. REV. 1027, 1050-55 (2012) (summarizing literature on disclosure from a variety of contexts); Cate, *supra* note 302, at 358-62 (citing studies on privacy notices); Edwards, *supra* note 302, at 229-33 (citing studies on disclosures mandated by the Truth in Lending Act).

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.[304] 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.

ii.   *Availability of comparable Internet service providers*

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 United States, this is a real problem. (See Box 10: Competitiveness of the Market for Broadband Internet Access in the United States below.) According to data published by the FCC in September 2014, 60% of housing units in the United States are in areas served by two wireline, facilities-based broadband Internet access providers, while 19% are in areas where only one such provider offers service.[305] This market structure has been characterized as "duopoly+/-."[306] 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

---

304. On this and the following, see the discussion in the text accompanying notes 251-78 above.

305. Fifteen percent 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 overbuilder). Six percent are in areas with no wireline provider. The FCC currently defines broadband as 4 Mbps download speed and 1 Mbps upload speed. *FCC Chairman Tom Wheeler: More Competition Needed in High-Speed Broadband Marketplace*, FCC (Sept. 4, 2014), https://apps.fcc.gov/edocs_public/attachmatch/DOC-329160A1.pdf [hereinafter *More Competition Needed*].

306. Farrell, *supra* note 211, at 201-02.

they would not be able to impose in a competitive market.[307] According to the FCC's National Broadband Plan, which was published in 2010, mobile Internet users in the United States 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.[308] In the European Union, 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.[309]

BOX 10

COMPETITIVENESS OF THE MARKET FOR BROADBAND INTERNET ACCESS IN THE UNITED STATES

Opponents of network neutrality regulation usually have a more optimistic view of the actual amount of competition in the United States than the one taken by this Article.[310] In particular, they use older FCC data based on Internet service availability by zip code,[311] which overstates the amount of competition;[312] treat mobile broadband Internet service as a substitute for rather than a complement to wireline Internet services;[313] and ignore or downplay the im-

---

307. *See, e.g.*, *id.* at 202-05.

308. FCC, CONNECTING AMERICA: THE NATIONAL BROADBAND PLAN 39-40 (2010), *available at* http://download.broadband.gov/plan/national-broadband-plan.pdf. *But see id.* at 39 ("These [numbers] likely overstate the coverage actually 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, for example, J. Scott Marcus, *Network Neutrality: The Roots of the Debate in the United States*, 43 INTERECONOMICS 30, 34-36 (2008).

309. Cave & Crocioni, *supra* note 184, at 58; *see also, e.g.*, *Impact Assessment*, *supra* note 293, at 91-93 (discussing the European Union's existing regulatory framework); Network Neutrality: Challenges and Responses in the EU and in the U.S., EUR. PARL. DOC. (IP/A/IMCO/ST/2011-02) 49-50 (2011), *available at* http://www.europarl.europa.eu /document/activities/cont/201108/20110825ATT25266/20110825ATT25266EN.pdf.

310. *See, e.g.*, Becker et al., *supra* note 45, at 502-06.

311. *See, e.g.*, *id.* at 503.

312. *See* U.S. GOV'T ACCOUNTABILITY OFFICE, GAO-06-426, TELECOMMUNICATIONS: BROADBAND DEPLOYMENT IS EXTENSIVE THROUGHOUT THE UNITED STATES, BUT IT IS DIFFICULT TO ASSESS THE EXTENT OF DEPLOYMENT GAPS IN RURAL AREAS 14-18 (2006), *available at* http://www.gao.gov/new.items/d06426.pdf.

313. *See, e.g.*, Becker et al., *supra* note 45, at 504-05 (treating mobile broadband Internet service as a substitute for wireline Internet service). *But see* Reply Comments of Free Press at 45-47, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Apr. 26, 2014) [hereinafter Free Press Open Internet Reply Comments], *available at* http://apps.fcc.gov/ecfs/document/view?id=7020437465 (arguing that mobile broadband Internet service is currently a complement to, *not* a substitute for, wireline Internet service); Comments of Free Press at 46 n.109, Inquiry Concerning the De-

pact of switching costs,[314] bundling, and differentiation in the market for Internet services on the effectiveness of competition.

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.[315] 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.[316] 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 (FTTP) 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 (FTTN) or on traditional digital subscriber line (DSL) networks are significantly lower.[317] In the United States, cable providers have generally upgraded their networks to DOCSIS 3.0, while DSL providers have been slow to upgrade their networks to FTTP.[318] As a result, 61% of homes in the United States have only one service provider—the cable provider—that can offer peak speeds of more than 50 Mbps down and 3 Mbps up. Only 16% have access to two such providers, and 21% do not have

---

ployment of Advanced Telecommunications Capability to All Americans in a Reasonable and Timely Fashion, GN Docket No. 09-137, A National Broadband Plan for Our Future, GN Docket No. 09-51 (Sept. 4, 2009), *available at* http://apps.fcc.gov/ecfs/document /view?id=7020037663 (same); Comments of Free Press at 40-44, 104-05, A National Broadband Plan for Our Future, GN Docket No. 09-51 (June 8, 2009), *available at* http://apps.fcc .gov/ecfs/document/view?id=6520219926 (same); FCC, *supra* note 308, at 40-42 (treating substitutability of mobile broadband Internet service for wireline Internet service as open to debate).

   314. *See, e.g.*, Becker et al., *supra* note 45, at 502-03.

   315. The following discussion draws in part on VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 262.

   316. *See* FCC, *supra* note 308, at 39-40; *see also* Reply Comments of AT&T Services, Inc. at 60-63, Protecting and Promoting the Open Internet, GN Docket No. 14-28, Framework for Broadband Internet Services, GN Docket No. 10-127 (Sept. 15, 2014), *available at* http://apps.fcc.gov/ecfs/document/view?id=7522634753.

   317. *See* FCC, *supra* note 308, at 42; *see also* Susan P. Crawford, *The Communications Crisis in America*, 5 HARV. L. & POL'Y REV. 245, 247-48 (2011).

   318. Crawford, *supra* note 317, at 246-49.

access to such service at all.[319] 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.[320]

The trend towards bundling differentiates the market further, giving providers additional market power.[321] 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.[322] 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 content or exclusive devices that are valuable to the customer.[323] 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 did not offer or support the iPhone.[324] 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

---

319. A mere 1.6% have access to three or more such providers. *See More Competition Needed*, *supra* note 305, at 1.

320. Susan Crawford has called this "the looming cable monopoly." Susan P. Crawford, *The Looming Cable Monopoly*, YALE L. & POL'Y REV. INTER ALIA (June 1, 2010, 2:30 PM), http://ylpr.yale.edu/inter_alia/looming-cable-monopoly.

321. The following discussion draws heavily on VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 263.

322. In a survey of broadband users in the United States, the FCC found that thirty-nine percent 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." *Broadband Decisions: What Drives Consumers to Switch— or Stick with—Their Broadband Internet Provider* 3 (FCC, Working Paper, 2010) [hereinafter *Broadband Decisions*], *available at* http://hraunfoss.fcc.gov/edocs_public/attachmatch /DOC-303264A1.pdf.

323. For example, an empirical study of competition between cable television and direct broadcast satellite (DBS) multichannel services showed that, while customers generally tend to switch from cable to DBS when the quality-adjusted price of cable increases substantially, the exclusive availability on cable 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. Andrew Stewart Wise & Kiran Duwadi, *Competition Between Cable Television and Direct Broadcast Satellite: The Importance of Switching Costs and Regional Sports Networks*, 1 J. COMPETITION L. & ECON. 679, 695-702 (2005).

324. 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." Joann S. Lublin & Spencer E. Ante, *Inside Sprint's Bet on iPhone*, WALL ST. J. (Oct. 4, 2011), http:// online.wsj.com/article/SB10001424052970203405504576603053795839250.html (internal quotation marks omitted).

customers, which allows them to impose restrictions they would not be able to impose in a perfectly competitive market.[325]

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.[326] For example, in 2008 and 2009, all mobile providers in France and Germany contractually banned the use of Internet telephony applications over mobile Internet connections.[327] 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 users in the United States, showed that many providers were engaging in discriminatory traffic management practices that targeted peer-to-peer file-sharing applications.[328] 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.

iii.   *Switching costs*

Third, the market for Internet services is characterized by significant switching costs that reduce consumers' willingness to switch and limit the effectiveness of competition.[329] Switching costs are the costs a customer incurs when switching to a competitor.[330] Switching costs make consumers' demand

---

325. That product differentiation may provide sellers with some degree of market power is well established in the literature. *See, e.g.*, DENNIS W. CARLTON & JEFFREY M. PERLOFF, MODERN INDUSTRIAL ORGANIZATION 203-05 (4th ed. 2005).

326. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 259.

327. This statement is based on the author's review of the terms of service of mobile Internet service providers in France (November 3, 2008) and Germany (July 19, 2009).

328. For a summary of Internet service providers' responses in that proceeding, see PARSONS, *supra* note 100, at 23-31.

329. The following six paragraphs draw heavily on VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 261-64. For an attempt to calculate the costs of switching broadband providers in France, see Jackie Krafft & Evens Salies, *The Diffusion of ADSL and Costs of Switching Internet Providers in the Broadband Industry: Evidence from the French Case*, 37 RES. POL'Y 706, 715-17 (2008); and Jackie Krafft & Evens Salies, *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 327, 333-41 (Claude Ménard & Michel Ghertman eds., 2009). *See also* Joseph Cullen & Oleksandr Shcherbakov, Measuring Consumer Switching Costs in the Wireless Industry, Attachment to Reply Comments of the Open Internet Coalition, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Apr. 26, 2010), *available at* http://apps.fcc.gov/ecfs/document/view ?id=7020438577 (estimating the explicit and implicit switching costs in the U.S. wireless industry at approximately $230).

330. For an in-depth overview of the economic literature on switching costs, see Joseph Farrell & Paul Klemperer, *Coordination and Lock-in: Competition with Switching Costs and*

less elastic, enabling a provider to charge a higher price.[331] 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, 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.[332] When switching from a broadband-over-cable service to a DSL service, a consumer will be charged for installation and will have to buy a DSL modem and other new equipment.[333] If Internet service is bundled with television and telephony (as is common in the United States), cancellation of the Internet service portion of the bundle may result in a loss, or a reduction, of the bundle discount, and the loss of that discount may then be a significant ongoing financial cost for the consumer.[334]

Network Effects, *in* 3 HANDBOOK OF INDUSTRIAL ORGANIZATION 1967 (Mark Armstrong & Robert H. Porter eds., 2007). For a treatment of switching costs in the context of information goods, see CARL SHAPIRO & HAL R. VARIAN, INFORMATION RULES: A STRATEGIC GUIDE TO THE NETWORK ECONOMY chs. 5-6, at 103-72 (1999).

331. HAL R. VARIAN, INTERMEDIATE MICROECONOMICS: A MODERN APPROACH 604-05 (5th ed. 1999); *accord* Jerry A. Hausman et al., *Residential Demand for Broadband Telecommunications and Consumer Access to Unaffiliated Internet Content Providers*, 18 YALE J. ON REG. 129, 164 (2001).

332. For example, HearUsNow.org, a project of the Consumers Union, found that a number of the top broadband providers in the United States charged early termination fees. At the time of a survey conducted 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. *See* John Dunbar, *Pulling Plug on Net Service Not Easy*, USA TODAY (Apr. 9, 2007, 12:39 PM), http://www.usatoday.com/tech /techinvestor/industry/2007-04-09-net-fees_N.htm?POE=TECISVA; Bob Williams, *The Next Big Thing in Broadband: Early Termination Penalties*, HEARUSNOW.ORG (Apr. 9, 2007, 12:18 PM), http://hearusnow.org/posts/126-the_next_big_thing_in_broadband_early _termination_penalties. In a survey of broadband users in the United States, 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." *Broadband Decisions*, *supra* note 322, at 3.

333. In a survey of broadband users in the United States, the FCC found that fifty percent 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." *Broadband Decisions*, *supra* note 322, at 3.

334. The customer may switch his whole bundle to the new provider, but that creates other problems. For example, it may make the decision to switch more complex or result in the loss of the preferred service offering, for example in television or telephony. In a survey

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.[335] If she cannot install the access hardware and software herself (which takes time and expertise), she must stay at home for the installation.[336] 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 resume and her website, and bear the risk of missing e-mail messages sent to the old address.[337] 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 five 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.[338] Medium and large businesses that switch Internet service providers will

---

of broadband users in the United States, the FCC found that thirty-nine percent 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." *Id.*

335. 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 AOL "by either making the cancellation process so painful for the customers that they could not bear to continue, or by simply ignoring their requests." Randall Stross, *Why Time Warner Has Fallen in Love with AOL, Again*, N.Y. TIMES (Sept. 25, 2005), http://www.nytimes.com/2005/09/25/business/25digi.html.

336. In a survey of broadband users in the United States, the FCC found that forty-three percent 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." *Broadband Decisions*, *supra* note 322, at 3.

337. On the use of provider-specific e-mail addresses as a way to increase switching costs in Internet services, see SHAPIRO & VARIAN, *supra* note 330, at 109-10. 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 a phone number when he switches providers. *See, e.g.*, 47 U.S.C. § 251(b)(2) (2013); 47 C.F.R. pt. 52 subpt. C (2014). In 2007, the FCC asked for comments on a petition to require e-mail providers to forward e-mail to a new e-mail address for a limited time. *See* Petition for Rulemaking, E-mail Address Portability, RM No. 11391 (FCC July 20, 2007), *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=6519560444); *see also* Public Notice, Report No. 2832 (FCC Sept. 26, 2007), *available at* http://apps.fcc.gov/ecfs/document/view?id=6519739081 (asking for statements and reply comments regarding the petition for rulemaking "In the Matter of E-mail Address Portability"). But it did not take any further steps in this proceeding.

338. Stross, *supra* note 335. In a survey of broadband users in the United States, the FCC found that thirty-four percent 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." *Broadband Decisions*, *supra* note 322, at 3.

often need to renumber their networks, which is a "costly, tedious and error-prone process."[339]

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. In the European Union, for example, the Universal Service Directive allows Internet service customers to switch providers in response to a change in disclosed discriminatory practices without incurring early termination fees.[340]

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 United Kingdom's market for fixed-line 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.[341] 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, which may occur when it has a large number of features or when it is bundled with other products.[342] This suggests that customers in the market for Internet access 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

339. B. Carpenter & Y. Rekhter, Internet Eng'g Task Force, RFC 1900, Renumbering Needs Work 2 (Feb. 1996), https://www.ietf.org/rfc/rfc1900.txt; B. Carpenter et al., Internet Eng'g Task Force, RFC 5887, Renumbering Still Needs Work 3-5 (May 2010), https://tools.ietf.org/html/rfc5887.

340. Universal Service Directive, *supra* note 291, art. 20, § 2.

341. Chatura Ranaweera & Jaideep Prabhu, *The Influence of Satisfaction, Trust and Switching Barriers on Customer Retention in a Continuous Purchasing Setting*, 14 INT'L J. SERVICE INDUSTRY MGMT. 374, 380, 390 (2003).

342. Thomas A. Burnham et al., *Consumer Switching Costs: A Typology, Antecedents, and Consequences*, 31 J. ACAD. MARKETING SCI. 109, 115-19 (2003).

than rational choice theory would predict.[343] 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.[344] 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 existing Internet service provider excludes applications or content they would like to use.

b.   *Lessons from experience with disclosure-only network neutrality regimes*

In sum, even if there is competition in the market for Internet access services, disclosure cannot replace substantive regulation as a tool to discipline providers. The experience in Europe and Canada and in the market for mobile Internet services in the United States 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 United States.[345] The European legal framework does not prohibit restrictions on end users' use of applications or services, but it 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) showed, many Internet service customers in the European Union are subject to restrictions on their fixed or mobile Internet services.[346] A recent study showed

---

343. 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. William Samuelson & Richard Zeckhauser, *Status Quo Bias in Decision Making*, 1 J. RISK & UNCERTAINTY 7, 37-38 (1988). 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. Jack W. Brehm, *Postdecision Changes in the Desirability of Alternatives*, 52 J. ABNORMAL & SOC. PSYCHOL. 384, 389 (1956).

344. 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). *See* Stefano DellaVigna & Ulrike Malmendier, *Contract Design and Self-Control: Theory and Evidence*, 119 Q.J. ECON. 353, 381-93 (2004).

345. *See supra* notes 305-08 and accompanying text (describing the U.S. market); *supra* Box 10 (providing additional context about the U.S. market); *supra* note 309 and accompanying text (describing the European market).

346. *A View of Traffic Management and Other Practices Resulting in Restrictions to the Open Internet in Europe: Findings from BEREC's and the European Commission's Joint Investigation*, BoR (12) 30 (May 29, 2012) [hereinafter *BEREC View of Traffic Management*], *available at* http://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/45-berec-findings-on-traffic-management-pra_0.pdf.

widespread discriminatory network management in the United Kingdom.[347] In Canada, the 2009 investigation of the 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.[348]

Under the FCC's Open Internet Order, providers of mobile Internet services in the United States 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.[349] 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 United States is considerably more competitive than the market for wireline Internet services.[350] Examples are Verizon Wireless's conduct towards tethering applications;[351] Verizon Wireless's, AT&T's, and T-Mobile's actions towards Google Wallet;[352] and AT&T's actions towards FaceTime.[353] 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 United States—competition does not prevent Internet service

---

347. Cooper, Analysis, *supra* note 297; Cooper, Thesis, *supra* note 297, ch. 6, at 131-70.

348. For an overview of Canadian providers' network management practices as disclosed during the proceeding, see PARSONS, *supra* note 100, at 15-31. 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. *See* Sarah Schmidt, *Complaints About Online Traffic Delays Accelerating, Says CRTC*, CANADA.COM (Jan. 12, 2012), http://www.canada.com/life/Complaints+about+online +traffic+delays+accelerating+says+CRTC/5986923/story.html; *see also* Geist, *supra* note 271 (describing Rogers's traffic-management practices).

349. 47 C.F.R. § 8.3 (2014); *id.* § 8.5(a), *invalidated by* Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014); *see also Open Internet Order*, 25 FCC Rcd. 17,905, 17,938-39 (2010) (report and order) (describing the obligation to disclose "[a]pplication-[s]pecific [b]ehavior" under 47 C.F.R. § 8.3 (italics omitted)), *vacated in part*, *Verizon*, 740 F.3d 623.

350. *See supra* notes 305-08 and accompanying text.

351. Barbara van Schewick, *Public Interest Requires Public Input: Verizon/Android Tethering*, INTERNET ARCHITECTURE & INNOVATION (June 30, 2011), https:// netarchitecture.org/2011/06/public-interest-requires-public-input-verizonandroid-tethering.

352. Barbara van Schewick, *Is Verizon Wireless Illegally Blocking Google Wallet? It's Time for the FCC to Investigate*, INTERNET ARCHITECTURE & INNOVATION (Dec. 19, 2011), https://netarchitecture.org/2011/12/is-verizon-wireless-illegally-blocking-google-wallet-its -time-for-the-fcc-to-investigate.

353. Cecilia Kang, *AT&T Faces Complaint over iPhone FaceTime Blocking*, WASH. POST POST TECH (Sept. 18, 2012, 9:08 AM ET), http://wapo.st/1yRD4ql; Chris Ziegler, *AT&T Only Allowing FaceTime over Cellular on Mobile Share Plans, No Extra Charge*, VERGE (Aug. 17, 2012, 4:29 PM), http://www.theverge.com/2012/8/17/3250228/att-facetime -over-cellular-ios-6-mobile-share.

providers from interfering with applications, content, or services on their networks, even if, as in the United States and the European Union, network providers are required to disclose any discriminatory conduct that occurs.[354]

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.[355] (See Box 11: The Benefits of Disclosure Rules.)

BOX 11
THE BENEFITS OF DISCLOSURE RULES[356]

While mandatory disclosure cannot replace substantive network neutrality rules, it is an important complement to such 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 United States compete based on maximum upload and download speeds and price. If, however, customers are unable to note the differences between the offerings along other dimensions (e.g., how oversubscribed the network is, how often traffic management is used, how traffic is prioritized), they cannot take these factors into account when making a decision, and network providers will not have an incentive to compete on these factors. Thus, disclosing these characteristics along with more detailed performance measures would not only help con-

---

354. 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* Barbara van Schewick, Comments on the European Commission's Public Consultation on Specific Aspects of Transparency, Traffic Management and Switching in an Open Internet at 19-21 (Oct. 15, 2012), *available at* http://ec.europa.eu/information_society/newsroom/cf/dae/document.cfm?doc_id=3168.

355. Network neutrality proponents generally support adopting disclosure rules as a complement to substantive regulations. *See, e.g.*, *id.*; *see also* Free Press Open Internet Reply Comments, *supra* note 313, at 17-21; Open Internet Coalition Comments, *supra* note 36, at 86-92; Public Interest Comments, *supra* note 36, at 63-67. In the United States, network providers have generally argued against any mandatory disclosure rules, whether as a substitute or a complement to substantive network neutrality regulation. *See, e.g.*, Comments of AT&T Inc. at 188-96, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter AT&T Open Internet Comments], http://fjallfoss.fcc.gov/ecfs/document/view?id=7020377217; Verizon Open Internet Comments, *supra* note 34, at 132.

356. The text in Box 11 draws directly on van Schewick, Official Testimony, *supra* note 68, at 1-2.

sumers make more informed choices, but also motivate Internet service providers to compete along these previously hidden dimensions.[357] 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. When asked by a reporter and later by the Electronic Frontier Foundation, Comcast denied that it was interfering with BitTorrent.[358] 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.[359]

2. *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

---

357. *See, e.g.*, Report of Bill St. Arnaud ¶¶ 40-42, 49-51, Initial Comments of Campaign for Democratic Media, Attachment C, Part 1, Review of the Internet Traffic Management Practices of Internet Service Providers, CRTC 2008-19 (Feb. 23, 2009) (Can.) [hereinafter St. Arnaud CRTC Report], *available at* http://www.crtc.gc.ca/public/partvii/2008/8646 /c12_200815400/1029987.zip (arguing in favor of requiring network providers to disclose oversubscription ratios); Benjamin Lennett, *Dis-Empowering Users vs. Maintaining Internet Freedom: Network Management and Quality of Service (QoS)*, 18 COMMLAW CONSPECTUS 97, 140-43 (2009) (same); *see also* COMM. ON COMMC'NS POLICY, IEEE-USA, NETWORK TRAFFIC MANAGEMENT AND THE EVOLVING INTERNET 12-18, 24 (2010) [hereinafter NETWORK TRAFFIC MANAGEMENT] (arguing in favor of standardization and disclosure of more detailed performance measures, including metrics for "bandwidth, latency, jitter, packet loss and availability"); IEEE-USA, POSITION STATEMENT: NETWORK TRAFFIC MANAGEMENT 1-3 (2010) [hereinafter IEEE POSITION STATEMENT] (same).

358. Marguerite Reardon, *Comcast Denies Monkeying with BitTorrent Traffic*, CNET (Aug. 21, 2007, 4:52 PM PDT), http://www.cnet.com/news/comcast-denies-monkeying-with -bittorrent-traffic; Seth Schoen, *Comcast and BitTorrent*, ELEC. FRONTIER FOUND. (Sept. 13, 2007), http://www.eff.org/deeplinks/2007/09/comcast-and-bittorrent. This sentence and the next sentence are adapted from VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 261.

359. 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. *See* ECKERSLEY ET AL., *supra* note 273, at 1-2; Peter Svensson, *Comcast Blocks Some Internet Traffic*, WASH. POST (Oct. 19, 2007, 6:32 PM), http://wapo.st/1Kx9KZM.

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.[360]

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.[361]

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.[362]

In response, the two nondiscrimination rules in this Subpart 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 Subpart 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, as shown in Table 1: Similarities and Differences Between the Approaches below.[363]

---

360. *See supra* notes 2, 61.

361. On these factors and their economic, social, cultural, and political impact, see Boxes 3 and 4 and notes 56-65 and accompanying text above.

362. See note 3 above for background on deep packet inspection in general, a specific example, and a discussion of the state of deep packet inspection deployment.

363. As the discussion will show, only the second approach meets this goal.

TABLE 1
Similarities and Differences Between 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, discussed in Part II.D.2.a, 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 13: Terminology: "Application" and "Class of Application" below.)

By contrast, the rule proposed by this Article—the second approach, discussed in Part II.D.2.b—is based on the insight that 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 application's characteristics.)[364] In line with this assessment, the proposed rule bans application-specific discrimination but allows application-agnostic discrimination.

The proposed rule 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 approach and the rule proposed by this Article differ in how they treat discrimination among classes of applications that are not alike, as shown in Table 1: Similarities and Differences Between the Approaches above.

---

364. For a more detailed discussion of the terms "application-specific," "discrimination based on application," and "discrimination based on class of application," see notes 384, 443-44, 449-50 and accompanying text below. *See also infra* Box 19.

The proposed rule 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).

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.

By contrast, requiring network providers to treat like traffic alike does not accurately distinguish between socially beneficial and socially harmful differential treatment and creates considerable social costs. Such a rule removes the application-agnosticism of the network, allows network providers to deliberately or inadvertently distort competition among applications or classes of applications, and violates the principles of user choice and innovation without permission that have fostered application innovation in the past. Due to the ambiguities surrounding the definition of "like," like treatment provides little certainty to the market and creates high costs of regulation.

a. *The first approach: 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.[365] (Again, I use "applications" as shorthand for "applications, content, services, and uses." The terms "application" and "class of application" are defined in Box 13: Terminology: "Application" and "Class of Application" below.) Thus, the approach requires network providers to treat like traffic alike. This requirement is often called "like treatment."[366] The nondiscrimination rule in the merger conditions

---

365. 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 Part II.D.2.b. *See also infra* Box 19.

366. On like treatment, see Wu, *supra* note 205, at 42-43. For criticism of allowing like treatment, see, for example, Free Press Open Internet Reply Comments, *supra* note 313, at 13-14, 101-04 (rejecting calls to allow Internet service providers to prioritize certain classes of traffic in the context of the reasonable network management exception); Center for Democracy & Technology Comments, *supra* note 36, at 29-30, 40 (arguing against allowing like treatment in the nondiscrimination rule and as reasonable network management); and M.

of the AT&T/BellSouth merger has been interpreted as requiring like treatment. (See Box 12: Like Treatment and the AT&T/BellSouth Merger Conditions 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 determination.[367] 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.[368]

CHRIS RILEY & ROBB TOPOLSKI, THE HIDDEN HARMS OF APPLICATION BIAS (2009). For support of allowing like treatment, see, for example, AT&T Open Internet Comments, *supra* note 355, at 187-88; Comments of Covad Communications Co. at 7-8, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010), *available at* http://apps.fcc.gov/ecfs/document/view?id=7020373797; and Comments of Cox Communications, Inc. at 23-30, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010) [hereinafter Cox Comments], *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=7020378714. *See also* Scott Jordan & Arijit Ghosh, *A Framework for Classification of Traffic Management Practices as Reasonable or Unreasonable*, 10 ACM TRANSACTIONS ON INTERNET TECH. 12:1, :10-11, :15, :19-20 (2010) (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").

367. See the rules discussed in Box 12 and note 383 below. For an exception in the context of the reasonable network management exception, see the Verizon-Google 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." Verizon & Google, *supra* note 218, at 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, *supra* note 205, at 42-43; Tim Wu, *The AT&T Network Neutrality Agreement*, WHAT'S NEW WITH WU (Dec. 29, 2006, 2:18 AM), https://web.archive.org/web/20070115072424/http://www.timwu.org/log/archives/81 (accessed via the Internet Archive index) [hereinafter Wu, *Neutrality Agreement*]. 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.*, PETER VAN DEN BOSSCHE, THE LAW AND POLICY OF THE WORLD TRADE ORGANIZATION: TEXT, CASES AND MATERIALS 320-400 (2d ed. 2008); Julia Ya Qin, *Defining Nondiscrimination Under the Law of the World Trade Organization*, 23 B.U. INT'L L.J. 215 (2005).

368. *See infra* Parts II.D.2.a.ii.B, II.D.2.a.ii.E.

BOX 12

LIKE TREATMENT AND THE AT&T/BELLSOUTH MERGER CONDITIONS

The nondiscrimination rule in the AT&T/BellSouth merger conditions, which prohibited AT&T/BellSouth from "privileg[ing], degrad[ing] or prioritiz[ing] any packet transmitted over AT&T/BellSouth's wireline broadband Internet access service based on its source, ownership or destination,"[369] has been interpreted as requiring like treatment. According to Tim Wu,

> While 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 . . . .[370]

Not all observers agree with this characterization of the agreement, though.[371]

BOX 13

TERMINOLOGY: "APPLICATION" AND "CLASS OF APPLICATION"

In this Article, the term "application" refers to a specific instance of a specific type of application. For example, Vonage is an application, as are Skype and 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

---

369. Merger Commitments at 8, Attachment to AT&T's Notice of Ex Parte Communication, Review of AT&T Inc. & BellSouth Corp. Application for Consent to Transfer of Control, WC Docket No. 06-74 (Dec. 28, 2006), *available at* http://apps.fcc.gov/ecfs/document/view?id=6518716381 ("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." (footnote omitted)).

370. Wu, *Neutrality Agreement*, *supra* note 367.

371. Comments of the National Cable & Telecommunications Ass'n, *supra* note 91, at 17-19 (disagreeing with Wu's interpretation of the merger conditions as allowing like treatment).

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 nondiscrimination rule described in this Subpart 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 Article, 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 Article, 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 Article, the term "class of application" is broader than that. In this Article, the group of applications that have the same application type (e.g., all e-mail applications, or "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."

i.   *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 the Xfinity TV website, Comcast's own Internet streaming video offering, differently from video streaming from other providers such as Netflix or YouTube.[372] Similarly, Comcast would not be allowed to count traffic from other providers' streaming video applications towards any monthly bandwidth usage cap while exempting traffic from the Xfinity TV website from the cap.[373] This would be discrimination among like applications.[374]

---

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

373. As explained above, the nondiscrimination rules discussed in this Article apply to all forms of differential treatment that make some applications or classes of applications relatively more attractive, not just to differential handling of packets in the network. *See supra* notes 81-90 and accompanying text.

In the past, Comcast's Internet service offerings had a monthly usage cap of 250 GB. At the time, data traffic from XfinityTV.com, Comcast's online streaming video offering, and from other streaming video providers all counted towards that cap. However, data traffic generated by Comcast's Xfinity TV 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, did not count towards the monthly bandwidth cap, while video traffic

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 video streaming from the Xfinity TV website, 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,[375] 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 videoconferencing 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.[376]

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.[377] Thus, in cases in which a network provider has an incentive to block an application or class of applications—for example, 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[378]—it often has an incentive to reach the same result by treating the targeted applications relatively worse than others (either by treating the other ap-

---

generated by other apps (for example, the HBO GO app) for the Xbox did count towards the cap. That behavior would violate the rule described in the text. For sources and additional discussion, see Box 5 above and Box 15 below.

374. In these examples, the group of like applications is the group of all streaming video applications.

375. 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. *See* sources cited *supra* note 29; *see also supra* Box 2; *supra* notes 29-30 and accompanying text.

376. This concern is discussed by, for example, van Schewick, Innovation Opening Statement, *supra* note 55, at 6; Jordan & Ghosh, *supra* note 366, at 12:14, :20; van Schewick, Open Internet Opening Statement, *supra* note 29, at 3; and the sources cited in note 109 above.

377. See the discussion in the text surrounding notes 97-105 above.

378. On incentives to block, see, for example, van Schewick, Official Testimony, *supra* note 68, at 5-6; and VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 222-70 (discussing increasing profits, managing congestion, and blocking unwanted content).

plications better or the targeted applications worse). The rule is designed to prevent this.[379]

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 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.[380] 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.[381]

ii.   *Allowing discrimination among classes of applications that are not alike*

While banning discrimination among like applications or classes of applications, this approach allows 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, the approach 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 discrimina-

---

379. This Article assumes that the case for a rule against blocking has been made. *See supra* note 28 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 nondiscrimination rule as well.

380. For a more detailed explanation, see notes 236-37 above.

381. For a more detailed explanation, see notes 51-60 above. *See also supra* Box 3.

tion between 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 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.[382]

The rule is agnostic as to who controls (or makes the decision regarding) the actual provision of the different types of service. (See Box 14: Like Treatment and Control over the Provision of Quality of Service below.) 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 did not discriminate among applications of that type.[383]

BOX 14

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 *partly provider-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 Communications described below in Box 17: Defining "Like" Based on Use,

---

382. 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 supra* note 367 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.*, VAN DEN BOSSCHE, *supra* note 367, at 320-400; Qin, *supra* note 367. As will be set out in more detail below, the interpretation in the text is not the only possible interpretation. This ambiguity of the term "like" is one of the key problems of this rule. *See infra* Parts II.D.2.a.ii.B, II.D.2.a.ii.E.

383. *See* Internet Freedom and Nondiscrimination Act of 2006, H.R. 5417, 109th Cong. § 3 (2006) ("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, S. 215, 110th Cong. § 2 (2007) ("[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 . . . .").

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.[384]

In a *fully 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 below in Box 16: Defining "Like" Based on Application Requirements is an example of a fully provider-controlled approach. Given the concerns described in Box 16, it is questionable whether that trial correctly classified certain time-sensitive applications. Thus, the Cox system may not comply with the rule described in the text, regardless of whether defining "like" based on application requirements is generally an acceptable interpretation of "like."

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 actu-

---

384. 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 nondiscrimination rule described in text would have to be decided in future adjudications. This question is beyond the scope of this Article. *See supra* notes 367-68; *infra* Parts II.D.2.a.ii.B, II.D.2.a.ii.E; *infra* Box 17.

ally get the type of service they need. Like treatment also creates considerable uncertainty and high costs of regulation.

A. *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[385]: 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 at which it negatively affects even applications that generally work well with best-effort service.[386] 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 anyway. 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.[387]

These cases, however, are only a nonrepresentative 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

---

385. *See supra* note 15 and accompanying text.

386. 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, *supra* note 15, at 494-99.

387. 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 infra* Part II.D.2.a.ii.B.

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.[388] The same incentive applies to Internet video offerings that may motivate users to cancel their subscription to the network provider's traditional video programming services,[389] or to Internet messaging applications that threaten cellular carriers' revenue from traditional text messaging services.[390] 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."[391] Since the nondiscrimination 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 15: 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 nondiscrimination rule.

---

388. *See, e.g.*, VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 240-43; *see also* Letter from AT&T to Ruth Milkman, Chief, Wireless Telecomms. Bureau, FCC 6-7 (Aug. 21, 2009), *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id =7020036306 (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").

389. *See, e.g.*, Competitive Impact Statement, *supra* note 258, at 11, 14-20, 37-39 (citing "[m]any internal documents" showing that Comcast views online video distributors as a competitive threat to its traditional cable video distribution offerings and describing Comcast's incentives to discriminate against unaffiliated online video providers).

390. *See* Brian X. Chen, *AT&T Chief Regrets Offering Unlimited Data for iPhone*, N.Y. TIMES BITS (May 4, 2012, 4:25 PM), http://bits.blogs.nytimes.com/2012/05/04/att-randall -stephenson (discussing the reduction in mobile carriers' text messaging revenues resulting from instant messaging applications); *Social Messaging Apps 'Lost Networks $13.9bn in 2011,'* BBC NEWS (Feb. 21, 2012, 7:47 PM ET), http://www.bbc.co.uk/news/technology -17111044 (same); Daniel Thomas, *Texting Revenues Hit by Web Services*, FIN. TIMES (Feb. 21, 2012, 8:19 PM), http://www.ft.com/intl/cms/s/0/5599fa26-5cb5-11e1-8f1f-00144feabdc0 .html (same).

391. Whether this incentive exists depends on how specialized services are regulated. *See, e.g.*, James B. Speta, *Supervising Managed Services*, 60 DUKE L.J. 1715, 1721-32, 1749-59 (2011).

BOX 15

EXAMPLES OF DIFFERENTIAL IMPACT: COMCAST'S DIGITAL VOICE SERVICE
AND COMCAST'S XFINITY TV APP FOR THE XBOX

Since the nondiscrimination 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 order against Comcast regarding its treatment of BitTorrent 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.[392] 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 its 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.[393]

The 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

---

392. 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 C. Bastian et al., Internet Eng'g Task Force, RFC 6057, Comcast's Protocol-Agnostic Congestion Management System 23 (Dec. 2010), https://tools.ietf.org/html/rfc6057.

393. *See* Letter from Dana R. Shaffer, Chief, Wireline Competition Bureau, & Matthew Berry, Gen. Counsel, FCC, to Kathryn A. Zachem, Vice President, Regulatory Affairs, Comcast Corp. (Jan. 18, 2009), *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=6520213534; Letter from Kathryn A. Zachem, Vice President, Regulatory Affairs, Comcast Corp., to Dana Shaffer, Chief, Wireline Competition Bureau, & Matthew Berry, Gen. Counsel, FCC 2 (Jan. 30, 2009) [hereinafter Zachem Letter], *available at* http://fjallfoss.fcc.gov/ecfs2/document/view.action?id=6520212400. For a brief description of Comcast's network management system, see Lennett, *supra* note 357, at 119-20.

cable service, and Microsoft's Xbox Live Gold subscription service to watch selected video-on-demand content from Comcast on the Xbox.[394] When Comcast introduced the app in the spring of 2012, Comcast's Internet service had a 250 GB monthly bandwidth cap.[395] Traffic associated with the Xfinity TV app to the Xbox did 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) did count towards the cap.[396] In general, this differential treatment of like applications would violate the nondiscrimination rule described in this Subpart.[397] Comcast, however, claimed that in contrast 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.[398] If this argument is correct,[399] 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.[400] 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.[401] In the United States, Comcast, RCN, and, most likely, Cox for a while

---

394. *See supra* Box 5.

395. In May 2012, Comcast suspended enforcement of the 250 GB monthly bandwidth cap and started trialing several different data usage management approaches in selected markets. *See supra* note 83.

396. *FAQs: Xbox 360*, *supra* note 84.

397. *See supra* Box 5; *supra* notes 81-82, 373 and accompanying text.

398. Werner, *supra* note 86.

399. Whether the Open Internet Rules apply to the Xfinity TV app is outside the scope of this Article.

400. van Schewick, Official Testimony, *supra* note 68, at 5-6; VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 264-66.

401. *See* PARSONS, *supra* note 100, at 23-31 (summarizing the filings by Canadian Internet service providers describing their network management practices). 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

managed traffic on their networks by selectively interfering with BitTorrent and other peer-to-peer file-sharing applications but not with other applications.[402] In 2009, BT throttled streaming video of users subscribing to its "Up to 8 Mbps Option 1" broadband plan to 896 kilobits per second between 5:00 PM and midnight.[403] And according to Neelie Kroes, who at the time was Vice President of the European Commission responsible for the Digital Agenda, data published by BEREC in June 2012 show that around twenty percent of fixed Internet service providers (spread across virtually all European Union member states) impose restrictions on peer-to-peer file-sharing applications during peak times. These restrictions can affect up to ninety-five percent of users in a country.[404]

When a network provider singles out a class of like applications for special treatment without discriminating among 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 applications that affects user behavior and the applications' chances in the marketplace. As the cofounders 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

that had not announced an intention to phase out that policy. *See* Geist, *supra* note 271; Schmidt, *supra* note 348.

402. RCN Letter, *supra* note 103, at 2, 4; Comcast Description, *supra* note 102, at 1, 9. Cox seems to have actively managed peer-to-peer file sharing in 2008 as well. Marcel Dischinger et al., *Detecting BitTorrent Blocking*, 2008 PROC. 8TH ACM SIGCOMM CONF. ON INTERNET MEASUREMENT 3, 7-8 (finding evidence of BitTorrent blocking by Comcast and Cox); Susan Davis, *Cox About to Feel Wrath of Net Neutrality Activists*, WALL ST. J. WASH. WIRE (May 15, 2008, 5:44 PM ET), http://blogs.wsj.com/washwire/2008/05/15/cox-about-to -feel-wrath-of-net-neutrality-activists (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" (internal quotation marks omitted)).

403. Cellan-Jones, *supra* note 101; *see also* Cooper, Analysis, *supra* note 297, at 21-22; Cooper, Thesis, *supra* note 297, chs. 5-7, at 105-211 (documenting widespread discriminatory network management in the United Kingdom).

404. The text closely paraphrases European Commissioner Neelie Kroes's description of these findings in her blog. Neelie Kroes, *Next Steps on Net Neutrality—Making Sure You Get Champagne Service if That's What You're Paying for*, EUR. COMMISSION NEELIE KROES (May 29, 2012), http://blogs.ec.europa.eu/neelie-kroes/netneutrality. For the detailed findings, see *BEREC View of Traffic Management*, *supra* note 346.

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.[405]

Differential treatment that makes a class of applications 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 gen-
erally or, when the differential treatment is used for congestion management,
during peak times when people want to use the Internet most.[406] Thus, treating
classes of applications differently may harm users and applications even if the
classes of applications are not alike.

B.   *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.[407]

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 differen-
tial treatment.[408] 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 discre-
tion when defining classes of applications or determining which class a specific
application should be assigned to. This allows them to deliberately or inadvert-
ently distort competition among applications or classes of applications.[409]

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 16: Defining "Like" Based on Application Require-
ments below.) Another approach may focus on whether the applications are
used for similar goals or whether they compete with each other. (See Box 17:
Defining "Like" Based on Use below.) Depending on which option is chosen, a

---

    405. Zediva Ex Parte Letter, *supra* note 236, at 3-4. As the Zediva letter and conversa-
tions with entrepreneurs and investors show, this is not a theoretical concern. For another
publicly documented example, see van Schewick, Oral Testimony, *supra* note 236, at 2.

    406. Barbara van Schewick, *The FCC's Open Internet Proposal—Lessons from Silicon
Valley*, INTERNET ARCHITECTURE & INNOVATION (Dec. 13, 2010), http://netarchitecture.org
/2010/12/the-fccs-open-internet-proposal-lessons-from-silicon-valley.

    407. van Schewick, December 2010 Ex Parte Letter, *supra* note 52, at 11-13.

    408. Center for Democracy & Technology Comments, *supra* note 36, at 29.

    409. van Schewick, December 2010 Ex Parte Letter, *supra* note 52, at 11-12.

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 applications. Sometimes, this may happen deliberately; sometimes, it may happen inadvertently.

BOX 16

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 Verizon-Google 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."[410] In 2009, Cox Communications, a cable provider in the United States, 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 deprioritized 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."[411] For example, web, Internet telephony, e-mail, or streaming video were classified as time-sensitive, while file access, software updates, or peer-to-peer protocols were classified as non-time-sensitive.[412]

There are reasons to believe that this system would have violated the non-discrimination rule described in this Subpart: 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 non-time-sensitive and, therefore, differently from other time-sensitive traffic would have violated the requirement to treat like traffic alike.[413]

---

410. Verizon & Google, *supra* note 218, at 1.

411. Chloe Albanesius, *Cox Says It Will Delay P2P Traffic, Software Updates*, PC MAG. (Jan. 28, 2009, 1:27 PM EST), www.pcmag.com/article2/0,2817,2339756,00.asp.

412. Cox Comments, *supra* note 366, at 24-30, app. A; Albanesius, *supra* note 411; *see also* M. CHRIS RILEY & BEN SCOTT, FREE PRESS, DEEP PACKET INSPECTION: THE END OF THE INTERNET AS WE KNOW IT? 6-8 (2009), *available at* http://www.freepress.net/files/Deep _Packet_Inspection_The_End_of_the_Internet_As_We_Know_It.pdf (criticizing Cox's approach); *infra* note 423.

413. *See infra* notes 418-23 and accompanying text.

BOX 17

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 ten dollars per month. This option provided enhanced Quality of Service only to Internet telephony applications but not to other latency-sensitive applications.[414] Thus, Shaw defined the class of applications to which it offered an enhanced type of service based on the use of the application ("Internet telephony"), not on the application's technical needs.

Usually, proponents of nondiscrimination rules that require like treatment do not specify how regulators should determine whether applications are alike. Thus, it is not clear whether regulators would endorse a definition of like based on the functionality provided by the application.[415] Such a definition would allow Internet service providers to discriminate among classes of applications with similar technical requirements (e.g., by providing low-delay service only to online gaming but not to Internet telephony),[416] enabling them to interfere with user choice and distort competition among classes of applications by steering users towards or away from certain classes of applications.[417] Thus, a definition of like based on the functionality provided by the application would allow Internet service providers to engage in exactly the kind of conduct—interfering with user choice and distorting competition among applications or classes of applications—that network neutrality rules are designed to prevent and should therefore be rejected.

Network providers may deliberately or inadvertently define classes in a way that hurts specific applications within a class. The CRTC's 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.[418] The Internet service providers argued that this was necessary to protect the performance of real-time applications (such as applications that stream video in

---

414. Russell Shaw, *Cable Broadband ISP's QoS Enhancement Surcharge Draws Vonage's Ire*, ZDNET (Mar. 7, 2006, 6:51 AM GMT), http://www.zdnet.com/blog/ip-telephony/cable-broadband-isps-qos-enhancement-surcharge-draws-vonages-ire/952.

415. *See supra* notes 367-68, 384 and accompanying text.

416. For a more detailed discussion of this example, see notes 424-27 and accompanying text below.

417. *See supra* notes 400-06 and accompanying text.

418. For an overview of the practices, see PARSONS, *supra* note 100, at 23-31.

real time) during times of congestion.[419] 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?[420] 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.[421] 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 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.[422] The concern that time-sensitive applications that use peer-to-peer file-sharing protocols may

---

419. *See, e.g.*, Bell Aliant Comments, *supra* note 224, ¶¶ 85-87.

420. Technically, YouTube and Vuze, like many other online video applications, use a technique called progressive download to create a near real-time streaming experience. *See* JAN OZER, VIDEO COMPRESSION FOR FLASH, APPLE DEVICES AND HTML5 (2011); *Vuze vs Cox the Copyright Cop*, P2PNET.NET (Feb. 2009), http://www.p2pnet.net/story/18402. *See generally* Andrew Odlyzko, The Delusions of Net Neutrality (Aug. 31, 2008) (unpublished manuscript), *available at* http://www.dtc.umn.edu/~odlyzko/doc/net.neutrality.delusions.pdf.

421. 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. *NBAR2 Protocol Library*, CISCO SYS. 11-14, http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6558/ps6616/product _bulletin_c25-627831.pdf (last updated Jan. 2012).

422. In the CRTC hearings that were part of that proceeding, some network provider representatives, when asked whether their traffic management system distinguished between time-sensitive and non-time-sensitive peer-to-peer applications, 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.*, *Review of the Internet Traffic Management Practices of Internet Service Providers: Proceedings on Telecom Public Notice CRTC 2008-19 Before the Can. Radio-Television and Telecomms. Comm'n* ¶¶ 4482-83 (2009) (presentation by Chris MacFarlane, Vice President, Corporate Engineering, Cogeco Cable Canada) (transcript available at http://www.crtc.gc.ca/eng /transcripts/2009/tt0710.htm); *id.* ¶¶ 5640, 5953-60 (presentation by Jean Brazeau, Senior Vice President, Regulatory Affairs, and Dennis Steiger, Group Vice President, Engineering, Shaw Communications Inc.) (transcript available at http://www.crtc.gc.ca/eng/transcripts /2009/tt0713.htm); *id.* ¶¶ 6253-56 (presentation by Jonathan Daniels, Vice President, Regulatory Law, Bell Canada) (transcript available at http://www.crtc.gc.ca/eng/transcripts/2009 /tt0714.htm). Unlike 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.

be harmed by practices that deprioritize 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.[423] 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.[424] 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.[425] By contrast, network providers seem to be less interested in provid-

---

The responses cited above seemed to imply that traffic management practices apply to all peer-to-peer file-sharing applications without distinguishing between peer-to-peer file-sharing applications that are time-sensitive and those that are not. *See id.* ¶¶ 3938-59 (presentation by Matt Stein, Vice President, Network Services, Primus Canada) (transcript available at http://www.crtc.gc.ca/eng/transcripts/2009/tt0709.htm) (explaining that peer-to-peer file-sharing is not time-sensitive).

423. For example, in 2009, Vuze's General Counsel, Jay Monahan, asked the FCC to investigate the impact of Cox's trial network management system on peer-to-peer traffic, expressing concern that it would hurt the performance of Vuze. Cox's trial system deprioritized 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. *See* Albanesius, *supra* note 411; *Vuze vs Cox the Copyright Cop, supra* note 420; *see also* Comments of Distributel Communications Ltd. Concerning the Part VII Applications by the Consumers Ass'n of Canada et al. (CAIP et al.) & by Vaxination Informatique (Vaxination) Requesting the CRTC to Review & Vary Telecom Decision CRTC 2008-108 ¶¶ 10-16, 40-46 (June 22, 2009) (Can.), *available at* http://crtc.gc.ca/public/partvii/2009 /8662/p8_200907727/1231235.pdf (asking the CRTC to review its decision regarding Bell Canada's throttling of its wholesale ADSL access service on the grounds that the CRTC based important parts of that decision on the assumption that all peer-to-peer file-sharing applications are not time-sensitive); Application by the Consumers' Ass'n of Canada et al. to Review & Vary Telecom Decision CRTC 2008-108 ¶¶ 126-30, 156-59, 166 (May 21, 2009) (Can.), *available at* http://crtc.gc.ca/public/partvii/2009/8662/p8_200907727/1140124.zip (same); Vaxination Informatique Application to Review & Vary Telecom Decision CRTC 2008-108 ¶¶ 95, 141-42 (May 20, 2009) (Can.), *available at* http://crtc.gc.ca/public/partvii /2009/8662/v42_200907826/1148017.zip (same). But the Commission rejected the applications for review and reiterated the claim that peer-to-peer file-sharing applications are not time-sensitive. *See* Canadian Association of Internet Providers et al. & Vaxination Informatique—Application to Review & Vary Certain Determinations in Telecom Decision 2008-108 Related to Bell Canada's Internet Traffic Management Practices, Telecom Decision CRTC 2009-677 ¶¶ 8-10 (Oct. 29, 2009) (Can.), *available at* http://crtc.gc.ca/eng /archive/2009/2009-677.htm.

424. "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.

425. *See, e.g.*, SANDVINE, TURNING GAMING INTO REVENUE WITH PACKETCABLE[TM] MULTIMEDIA 1 (2005) (on file with author) ("Once aware of the amounts and types of gam-

ing low-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.[426] 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.[427] 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 those applications used peer-to-peer file-sharing protocols,[428] but they may not have realized that there are applications, such as 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

---

ing 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 online gaming drives demand for exactly that type of service. In this way, an MSO [multiple-system operator, an industry term for cable operators] can profit from an otherwise peripheral market trend.").

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 an 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 associated with each application through the transport price alone. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 273-75. 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.

426. This observation is based on conversations with network provider employees.

427. For a real-world example of such an offering, see *Success Story: Service Innovation with Third Party Partnerships*, SANDVINE (Aug. 16, 2013), https://www.sandvine .com/downloads/general/success-stories/success-story-vox-telecom-service-innovation-with -third-party-partnerships.pdf (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).

428. See the responses of network provider representatives discussed above in note 422.

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, thus creating high costs of regulation.

C. *User choice*

Apart from removing the application-agnosticism of the network, "like treatment" also violates the principles of user choice and innovation without permission. Like application-agnosticism, these principles have been central to the Internet's ability 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.[429]

Under "like treatment," network providers, not users, choose which application should get which Quality of Service or differential treatment, thus violating the principle of user choice.[430] As I have explained elsewhere, the incentives of network providers and users are not necessarily aligned.[431] Network providers' incentive to offer low-delay service only to online gaming, but not to Internet telephony, or to reduce the performance of applications that may reduce their revenue from applications that are offered and provided separately from Internet access are examples of this phenomenon.[432] Thus, network providers do not always want to meet users' preferences. But even when they do, they may not be able to do so.[433] For example, if a network provider decides whether and when to offer Quality of Service, 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.[434] A user's desire

---

429. *See supra* Part I.

430. van Schewick, December 2010 Ex Parte Letter, *supra* note 52, at 12.

431. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 350-51.

432. *See supra* notes 424-27 and accompanying text; *see also* 388-91 and accompanying text.

433. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 351.

434. van Schewick, Official Testimony, *supra* note 68, at 7; *see also* Free Press Open Internet Comments, *supra* note 3, at 102-03; RILEY & SCOTT, *supra* note 412, at 8; Lennett, *supra* note 357, at 143-45; Yiannis Yiakoumis et al., *Putting Home Users in Charge of Their Network*, 2012 PROC. 2012 ACM CONF. ON UBIQUITOUS COMPUTING 1114, 1115; B. Briscoe

for Quality of Service 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.[435] Thus, any Quality of Service system that lets network providers determine whether and when to provide Quality of Service 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.[436]

D. *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.[437] In order to get Quality of Service, 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, thus violating the principle of innovation without permission.[438] This introduces considerable transaction costs. Certain types of innovators (e.g., innovators that develop an application at home in their free time, noncommercial innovators, or start-ups) may not have the resources necessary to engage in this type of negotiation with a potentially large number of network providers.[439] In addition, even if an innovator manages to contact a network provider, the innovator may not receive the appropriate Quality of Service for its application if the innovator fails to convince the network provider. This is an example of the more general phenomenon that requir-

---

et al., Internet Eng'g Task Force, Internet-Draft, Problem Statement: Transport Protocols Don't Have to Do Fairness 13-14 (July 14, 2008), https://tools.ietf.org/html/draft-briscoe -tsvwg-relax-fairness-01; *supra* note 147 and accompanying text; *infra* notes 484-85 and accompanying text.

435. This example is taken from Yiakoumis et al., *supra* note 434, at 1115.

436. *See, e.g.*, RILEY & SCOTT, *supra* note 412, at 8; Lennett, *supra* note 357, at 143-44.

437. van Schewick, December 2010 Ex Parte Letter, *supra* note 52, at 12-13.

438. Center for Democracy & Technology Comments, *supra* note 36, at 29.

439. Throughout the history of the Internet, many important innovations (including eBay, Facebook, Yahoo, Google, Apache Web Server, the World Wide Web, Flickr, and Blogger) have been developed by innovators of this type. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 204-13, 310-14, 318-28, 334-45 (discussing the importance of different types of low-cost innovators and including many examples); van Schewick, Open Internet Opening Statement, *supra* note 29, at 3-5.

ing cooperation or support from the network provider reduces the likelihood that innovative applications can be realized or successfully deployed.[440] 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.

E. *Certainty and costs of regulation*

In general, this rule—ban discrimination among like applications and classes of applications, but allow discrimination among classes of applications that are not alike—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 from offering 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.[441] If adjudicators clarify the interpretation of "like" in the context of individual adjudications, this uncertainty may be reduced over time.[442] 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 there-

---

440. *See supra* note 57 and accompanying text.

441. On the importance of certainty for network providers and application developers, see notes 234-37 and accompanying text above.

442. 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 supra* notes 238-40 and accompanying text.

fore be allowed. 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. 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.

    b.   *The best approach: ban application-specific discrimination, but allow application-agnostic discrimination*

Instead, regulators or legislators should adopt a nondiscrimination rule that clearly bans application-specific discrimination, but allows application-agnostic discrimination.[443] (Again, I use "applications" as shorthand for "applications,

---

    443. *See* van Schewick, December 2010 Ex Parte Letter, *supra* note 52, at 13-16; van Schewick, August 2010 Attachment, *supra* note 52, at 6-8; *see also* VAN SCHEWICK, *supra* note 79, at 52. In the Open Internet proceeding, this proposal was supported by, for example, networking experts, *e.g.*, Ex Parte Letter of NYSERNet, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 13, 2010) [hereinafter NYSERNet Ex Parte Letter], *available at* http://apps.fcc.gov/ecfs/document /view?id=7020923371; Reed, *supra* note 2; venture capitalists, *e.g.*, John Borthwick, *Neutrality or Bust*, TECHCRUNCH (Dec. 19, 2010), http://techcrunch.com/2010/12/19/neutrality; Brad Burnham, *Internet Access Should Be Application-Agnostic*, HUFFINGTON POST (Dec. 20, 2010, 9:44 AM ET), http://www.huffingtonpost.com/brad-burnham/internet-access -should-be_b_799028.html; *Regulation Strangulation*, AVC (Aug. 12, 2010), http://www .avc.com/a_vc/2010/08/regulation-strangulation.html; entrepreneurs, *e.g.*, Zediva Ex Parte Letter, *supra* note 236; and nonprofit organizations, *e.g.*, Ex Parte Letter of Council of Scientific Society Presidents, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 13, 2010) [hereinafter Council of Scientific Society Presidents Ex Parte Letter], *available at* http://apps.fcc.gov/ecfs/document /view?id=7020923181; Ex Parte Letter of North American Benthological Society, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 13, 2010) [hereinafter Benthological Society Ex Parte Letter], *available at* http://apps.fcc.gov/ecfs/document/view?id=7020923349; Ex Parte Letter of Botanical Society of America, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Dec. 8, 2010) [hereinafter Botanical Society Ex Parte Letter], *available at* http://apps.fcc.gov/ecfs/document/view?id=7020922879.

    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. For example, the Open Internet Coalition stated that to qualify as reasonable network management, the practice must be narrowly tailored to address a legitimate network management purpose. Open Internet Coalition Comments, *supra* note 36, at 15-17, 49-50. In addition to other criteria, the practice must "result[] in as little discrimination or preference as reasonably possible." *Id.* at 49. The Cen-

content, services, and uses.") Discrimination is application-specific if it is based on a particular application or class of applications, or, in other words, if it is based on criteria that depend on an application's characteristics ("application-specific criteria").[444] Application-specific criteria include what this Article calls "application"—the 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., Session Initiation Protocol (SIP) vs. Skype's proprietary protocol, or TCP vs. User Datagram Protocol (UDP)), or the application's technical requirements (e.g., latency-sensitive vs. non-latency-sensitive applications). (See also Box 19: A Technical Perspective on Application-Specific vs. Application-Agnostic Discrimination below.) Since the term "applications" stands for applications, content, services, and uses, the ban on application-specific discrimination ap-

---

ter for Media Justice and other public interest commenters similarly agreed that to qualify as reasonable network management, the practice must be narrowly tailored to address a legitimate network management purpose. Public Interest Comments, *supra* note 36, at 31-32, 35-41. These two proposals would allow application-agnostic differential treatment only as long as it is 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.

444. 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 Article, "application" refers to a specific instance of a specific type of application. *See supra* Box 13. 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 id.* 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 Session Initiation Protocol (SIP), all applications that use Hypertext Transfer Protocol (HTTP)) or transport-layer protocol (e.g., all applications that use TCP, all applications that use User Datagram Protocol (UDP)), or all applications that have similar technical requirements (e.g., all latency-sensitive applications, all latency-insensitive applications). A network provider discriminates "based 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.

plies equally to discrimination based on criteria that depend on characteristics of content or of a service or use. Thus, discrimination against certain content based on, for example, publisher, author, content type, subject matter, or viewpoint would also all be prohibited by this rule.

The rule should be coupled with an exception for reasonable network management that 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 18: The Exception for Reasonable Network Management below.)

This rule plays an important role in the FCC's Open Internet Order. The Order's nondiscrimination rule for fixed broadband access banned discrimination that is unreasonable. Whether discriminatory behavior complies with the rule proposed by this Article and described in this Subpart (i.e., whether it is application-agnostic) is one of the factors the FCC proposed to use to determine the reasonableness of discriminatory conduct under the Open Internet Order's nondiscrimination rule and exception for reasonable network management.

BOX 18

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 nondiscrimination 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 that 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.[445]

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 pur-

---

445. *See, e.g.*, van Schewick, Notice of Ex Parte Meetings, *supra* note 53, at 5; van Schewick, Official Testimony, *supra* note 68, at 4-8; van Schewick, Oral Testimony, *supra* note 236; van Schewick, *supra* note 406; van Schewick, *supra* note 263. During the Open Internet proceeding, the proposed exception was supported by, for example, networking experts, *e.g.*, NYSERNet Ex Parte Letter, *supra* note 443, at 2; venture capitalists, *e.g.*, Borthwick, *supra* note 443; Burnham, *supra* note 443; *Regulation Strangulation, supra* note 443; entrepreneurs, *e.g.*, Zediva Ex Parte Letter, *supra* note 236; and nonprofit organizations, *e.g.*, Council of Scientific Society Presidents Ex Parte Letter, *supra* note 443; Benthological Society Ex Parte Letter, *supra* note 443; Botanical Society Ex Parte Letter, *supra* note 443.

pose is "to maintain, protect, and ensure the efficient operation of a net-work."[446] Network management includes, for example, managing congestion or protecting the security of a network.[447] 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.[448] The treatment of network management practices under the proposed rule is described in more detail in Part II.D.2.b.i.B below.

The rule described in this Subpart 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 nondiscrimination 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 nondiscrimination 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 based on application.[449] Nor would a network provider 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.[450] 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 Box 19: 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

---

446. Public Interest Comments, *supra* note 36, at 37.

447. On the definition of network management, see Center for Democracy & Technology Comments, *supra* note 36, at 41-43; and Public Interest Comments, *supra* note 36, at 37-41.

448. This formulation mirrors proposals by many network neutrality proponents. *See, e.g.*, Open Internet Coalition Comments, *supra* note 36, at 48-50; Public Interest Comments, *supra* note 36, at 35-41.

449. On the meaning of "discrimination based on application," see note 444 above.

450. On the meaning of "discrimination based on class of application," see note 444 above.

"Up to 3 Mbps" Internet service packet).[451] During times of congestion, a net-work 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.[452] 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.

BOX 19

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.[453] 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.[454] 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 (IP)—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

---

451. *See* Center for Democracy & Technology Comments, *supra* note 36, at 25-26.

452. *See* Reply Comments of the Center for Democracy & Technology at 19-20, 22, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Apr. 26, 2010) [hereinafter Reply Comments of the Center for Democracy & Technology], *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id=7020437353 (arguing in favor of a strict nondiscrimination rule but asking the FCC to clarify that "the nondiscrimination rule shall 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"); Center for Democracy & Technology Comments, *supra* note 36, at 25-26.

453. *See* Reed, *supra* note 2. David Reed is one of the network engineers involved in the design of the Internet's original architecture and one of the authors of a famous paper that first identified and described the "end-to-end arguments," a key design principle of the Internet. On the end-to-end arguments, see notes 2, 56, and 61 above.

454. For a detailed explanation and references to the literature, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 52, 56-57.

other words, in the Internet's original architecture, the layering principle forced the network to be application-blind and application-agnostic.[455]

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 the Internet Protocol on behalf of higher-layer protocols; nor does it prohibit network providers from making assumptions about the meaning of these higher-layer protocol data units. But if network providers have somehow acquired information about the content or meaning of the payload of the Internet Protocol 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 nondiscrimination rule creates the same effect that compliance with the layering principle would have created.[456] Translated into the less technical terminology used throughout this Article, 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

---

455. *See* Reed, *supra* note 2 (arguing that the Internet's original architecture forced the Internet to be 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. 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 was unable to make distinctions among applications. Thus, it was application-blind and, like all application-blind networks, necessarily application-agnostic. *See supra* Box 4.

456. *See* Reed, *supra* note 2. Reed argues that the rule proposed by this Article and described in this Subpart effectively requires by law what the original design of the Internet required by code. He then proposes to replace the proposed rule with one that "requir[es] those who offer Internet service to implement the Internet design as it was intended." *Id.* Reed goes on to say that,

> [i]n 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 [van Schewick] is correct that only application-agnostic discrimination makes sense.

Id.

principle and be application-blind. It only requires the network to be application-agnostic.[457] (For an explanation of the rationale behind this decision, see Box 4: Application-Agnostic vs. 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 Internet Protocol packets traveling through the network would not be application-agnostic.[458] 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.[459]

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 nondiscrimination rule proposed here does not prevent network providers from collecting application-specific information. It only prevents them from using this information to, for example, discriminate

457. 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 above.

458. *See* Reed, *supra* note 2.

459. *See also id.* (arguing that the rule proposed in the text effectively requires by law what the original design of the Internet required by code); *supra* Box 19.

among data packets or charge differently based on application-specific criteria.[460] 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.[461] (The rationale behind this decision is explained in Box 4: Application-Agnostic vs. Application-Blind above. See also Box 19: 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,[462] and to all application-specific discriminatory conduct, regardless of whether the conduct is capable of monopolizing the market for the affected applications.[463] 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.[464] 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.[465] 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 different types of service ("Quality of Service") 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.[466] 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

---

460. For a full discussion of the kind of differential treatment subject to the rule, see Part II.A above.

461. 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 above.

462. See the discussion of the disclosure rule in Part II.D.1 above. For a full analysis, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 255-64. *See also supra* notes 204-05.

463. *See supra* notes 188-203 and accompanying text.

464. *See* VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, chs. 6-9, at 215-375; *see also supra* notes 198-203, 372-81, 383-442 and accompanying text.

465. *See supra* notes 206-17 and accompanying text.

466. *See supra* notes 366-440 and accompanying text.

potential. Network providers can, for example, manage their networks in application-agnostic ways, price discriminate based on application-agnostic criteria, or differentiate their services by offering Quality of Service in line with the rule.[467] 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 nondiscrimination rule in narrowly tailored ways.[468]

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.[469] 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 exactly what users want.[470] Network providers' private interests and the public's 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.[471] 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 resulting from improved democratic discourse.[472] 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.[473]

Compared with other proposals for nondiscrimination rules, this 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.[474] 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

---

467. On network management, see Part II.D.2.b.i.B below. On Quality of Service, see Part II.D.2.b.i.A below.

468. *See supra* Box 18.

469. This Subpart summarizes arguments that I have developed in detail elsewhere. *See* sources cited *infra* notes 470-74.

470. *See supra* notes 431-36 and accompanying text.

471. For a detailed discussion, see VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 355-71 (describing the public interest); *id.* at 371-75 (describing network providers' private interests and why they diverge from the public interest).

472. *Id.* at 373-74; *see also* Frischmann, *supra* note 32, at 1009-12; Frischmann & van Schewick, *supra* note 36, at 400-03, 424-25; Hogendorn, *supra* note 205, 195-203.

473. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 374-75.

474. On these factors and their economic, social, cultural, and political impact, see notes 56-65 and accompanying text above. *See also supra* Boxes 3-4.

value of the Internet for users and for society, but also is 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.

i.   *Allowing the network to evolve*

The proposed rule does not constrain the evolution of the network infrastructure more than is necessary to reach the goals of network neutrality regulation. It provides room for networks to evolve.[475]

A. *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;[476] and (3) the network provider is allowed to charge only its own Internet service customers for the use of the different classes of service.[477]

---

475. For early versions of the arguments in this Subpart, see generally van Schewick, December 2010 Ex Parte Letter, *supra* note 52, at 14-15; van Schewick, August 2010 Attachment, *supra* note 52, at 7.

476. 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* Reply Comments of the Center for Democracy & Technology, *supra* note 452, at 20-21, 23 (stating that the 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"); Reply Comments of Google Inc. at 36, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Apr. 26, 2010), *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id =7020438889; Center for Democracy & Technology Comments, *supra* note 36, at 26-27; Free Press Open Internet Comments, *supra* note 3, at 103-04 (stating that if there is a demonstrated need for priority, users should make the choice); Open Internet Coalition Comments, *supra* note 36, at 50-51 (discussing some of the potential options for user-controlled Quality of Service); Comments of Skype Communications S.A.R.L. at 16-20, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (Jan. 14, 2010), *available at* http://fjallfoss.fcc.gov/ecfs/document/view?id =7020377906; van Schewick, Official Testimony, *supra* note 68, at 7-8 (discussing the question in the context of the reasonable network management exception); NETWORK TRAFFIC MANAGEMENT, *supra* note 357, at 16-18, 23 (supporting the model of user-controlled Quality of Service described in the text); IEEE POSITION STATEMENT, *supra* note 357, at 2-3; Jordan & Ghosh, *supra* note 366, at 12:14, :21-:22; Lennett, *supra* note 357, at 143-45.

477. This condition prohibits a provider of last-mile Internet access from charging the end user at the other end of the connection (e.g., an application provider) as well as from charging interconnecting networks, content delivery networks, or application providers that

For example, a network provider could offer a low-delay service, a best-effort service, a less-than-best-effort 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 might use the low-delay service for Internet telephony, another might use it for online gaming, and a third might use it for e-mail, if that is what that user wants. This type of user-controlled Quality of Service is technically feasible.[478] (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 Part II.D.2.b.iii below.)

While the first two conditions directly follow from the proposed nondiscrimination rule,[479] the third condition is based on additional considerations and would need to be encoded separately.[480]

---

directly interconnect with a last-mile Internet service provider. See Question 3 in Part III.D.2.b.iii below.

I explain the rationale for this criterion elsewhere. *See* sources cited *supra* note 29. 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 Article. For a short overview of the options, see Box 2 and notes 29-30 and accompanying text above.

478. This statement is based on many conversations with networking experts. *See also* Reply Comments of the Center for Democracy & Technology, *supra* note 452, at 20-21; Center for Democracy & Technology Comments, *supra* note 36, at 27; Jordan & Ghosh, *supra* note 366, at 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 Presentation, *supra* note 171, at 25, 29. For a prototype targeting home users, see Yiakoumis et al., *supra* note 434, at 1116-18, which won the Grand Prize at the Imagine App Challenge at the Cable Show 2012. *See* Mari Silbey, *Stanford Team Wins Cable Show App Challenge*, SMARTPLANET (May 24, 2012, 7:55 AM PDT), http://www.smartplanet.com/blog/thinking -tech/stanford-team-wins-cable-show-app-challenge/11743; Todd Spangler, *Cable Show 2012: Stanford Team Wins 'App Challenge' with Bandwidth-Priority System*, MULTI-CHANNEL NEWS (May 23, 2012), http://www.multichannel.com/news/mobile/cable-show -2012-stanford-team-wins-app-challenge-bandwidth-priority-system/306206. For a more detailed discussion, see the response to Question 9 in Part II.D.2.b.iii below.

The form of user-controlled Quality of Service described in the text does not violate the broad version of the end-to-end arguments. VAN SCHEWICK, ARCHITECTURE AND INNOVATION, *supra* note 2, at 106-07. On the broad version of the end-to-end arguments, see note 2 above.

479. 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).

A network provider that 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 is one of the main motivations behind proposals to ban Quality of Service.[481] 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.[482]

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 dependent not 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

---

480. Under the network neutrality regime that I propose, the restrictions on charging application providers that are not an Internet service provider's Internet access customers would be captured by a prohibition on all forms of access fees. I explain the rationale for a ban on access fees elsewhere. *See* sources cited *supra* notes 29-30. My network neutrality regime would also include a rule that prohibits last-mile Internet service providers from charging interconnecting networks, content delivery networks, or application providers that interconnect directly with last-mile Internet service providers for the termination and transport of their traffic to and from that Internet service provider's subscribers. On access fees, see note 62 (defining the term) above as well as Box 2 and notes 29-30 and accompanying text above. On interconnection, see the response to Question 3 in Part II.D.2.b.iii below. These questions are outside the scope of this Article.

481. *See supra* notes 110-11 and accompanying text.

482. 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 main text of this Article, 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 Article. For a discussion of the problem and of potential solutions, see van Schewick, Background Paper, *supra* note 29, at 10-11. The European Union has adopted a similar rule following its review of the regulatory framework for telecommunications services. *See* Directive 2009/136/EC, *supra* note 292, art. 22(3), at 25; *Impact Assessment*, *supra* note 293, at 92, 95-97, 101.

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.[483]

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.[484] As discussed repeatedly throughout this Article,[485] 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.

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.[486] 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

---

483. *See* Reply Comments of the Center for Democracy & Technology, *supra* note 452, at 27.

484. van Schewick, Official Testimony, *supra* note 68, at 7 (discussing the question in the context of the reasonable network management exception); RILEY & SCOTT, *supra* note 412, at 8; Lennett, *supra* note 357, at 143-45.

485. *See supra* notes 147, 430-36 and accompanying text.

486. *See* Center for Democracy & Technology Comments, *supra* note 36, at 26-27.

and intuitive.[487] For example, Skype could offer a simple button through which users could request high quality. The user would only need 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 could even offer different interfaces to support users with varying levels of sophistication.[488]

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 potentially operated by different users and running applications with diverse needs, have become quite complex, making them more difficult to manage and secure. Today, companies like Meraki, PowerCloud, and 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.[489] Managing the user's Quality of 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.[490]

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.

B. *Network management*

The proposed nondiscrimination rule allows any differential treatment that is application-agnostic. This includes any network management practices that treat traffic differently based on application-agnostic criteria.[491] In addition, the

---

487. 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., *supra* note 434.

488. *See infra* Part II.D.2.b.iii, Question 9.

489. *See, e.g.*, Nick Feamster, *Outsourcing Home Network Security*, 2010 PROC. 2010 ACM SIGCOMM WORKSHOP ON HOME NETWORKS 37 (proposing an approach for outsourcing home network security); Yiannis Yiakoumis et al., *Slicing Home Networks*, 2011 PROC. 2ND ACM SIGCOMM WORKSHOP ON HOME NETWORKS 1 (proposing a mechanism for outsourcing home network management more generally and describing an initial prototype that has been deployed).

490. *See* Yiakoumis et al., *supra* note 489, at 3 (explicitly discussing this possibility).

491. Network management practices that treat traffic differently based on application-agnostic criteria would be allowed under the proposed nondiscrimination 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 supra* note 443. On

proposed reasonable network management 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.[492]

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 the principle of user choice as much as possible.[493] 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),[494] network providers may limit the amount of capacity available to users of that link based on application-agnostic criteria. 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.[495] Even during times of congestion, network providers cannot, however, interfere with how users use the (limited) capacity available to them, for example, 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

---

the definition of network management, see Box 18 above. For a longer discussion of the policy arguments driving the treatment of network management measures proposed in the text, see van Schewick, Official Testimony, *supra* note 68, at 4-8; van Schewick Oral Testimony, *supra* note 236; and van Schewick, *supra* note 406. *See also supra* notes 213-16, 251-58, 268-75 and accompanying text.

492. For a more detailed description of the reasonable network management exception proposed by this Article, see Box 18 above.

493. Application-agnostic network management preserves application-agnosticism and user choice. These principles are compromised only if application-agnostic network management is impossible.

494. 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 supra* Box 7.

495. *See* Reply Comments of the Center for Democracy & Technology, *supra* note 452, at 19-20, 22 (arguing in favor of a strict nondiscrimination rule, but asking the FCC to clarify that "the nondiscrimination rule shall 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"); Center for Democracy & Technology Comments, *supra* note 36, at 25-26.

when a link's average utilization is high),[496] network providers could allow users to choose which applications to prioritize or otherwise treat differently during these times.[497] 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 management would be consistent with the nondiscrimination rule and reasonable network management exception proposed above.[498]

Tools for application-agnostic congestion management are available today. For example, Comcast, the largest provider of broadband Internet access services in the United States,[499] adopted an application-agnostic congestion management system in response to the FCC's order against Comcast in 2008.[500] 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 HSI [High Speed Internet] customers."[501] 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.[502] 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.[503] Network providers could

---

496. On the use of the term "congestion," see note 494 above.

497. van Schewick, Official Testimony, *supra* note 68, at 7-8; Jordan & Ghosh, *supra* note 366, at 12:17-:20.

498. See the discussion of Quality of Service in the previous Subpart.

499. Press Release, Leichtman Research Grp., *supra* note 194.

500. For descriptions of Comcast's application-agnostic network management system, see Comcast Corp. Description of Planned Network Management Practices to be Deployed Following the Termination of Current Practices, Letter from Kathryn Zachem, *supra* note 102, Attachment B; Zachem Letter, *supra* note 393, at 1-3; and Bastian et al., *supra* note 392.

501. Bastian et al., *supra* note 392, at 23.

502. *See Congestion Exposure (CONEX): Charter for Working Group*, INTERNET ENGINEERING TASK FORCE, https://datatracker.ietf.org/wg/conex/charter (last visited Jan. 7, 2015).

503. For an overview, see generally Arnaud Jacquet et al., *Policing Freedom to Use the Internet Resource Pool*, 2008 PROC. 2008 ACM CONEXT CONF.; and B. Briscoe et al., Internet Eng'g Task Force, RFC 6789, Congestion Exposure (ConEx) Concepts and Use Cases (Dec. 2012), https://tools.ietf.org/html/rfc6789.

use this information, for example, to allocate bandwidth among users during times of congestion based on their contribution to congestion, charge users based on their contribution to congestion, or 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 and users on the other hand.[504] Application-agnostic network management practices remove this incentive, freeing resources for 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.[505]

ii.   *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.[506] Contrary to those approaches, the rule is immediately applicable to all industry participants. This not only removes the need to readjudicate similar cases again and again;[507] it also avoids the intertemporal inconsistencies across

---

504. *See supra* notes 269-72 and accompanying text.

505. For a more detailed description of the reasonable network management exception proposed by this Article, see Box 18 above.

506. *See generally* 1 PIERCE, *supra* note 241, § 6.8, at 497-98, 500 (describing how rulemaking is more efficient and fairer than standards-based approaches). For a more detailed discussion, see Parts II.C.3 and II.C.4 above.

507. *See generally* 1 PIERCE, *supra* note 241, § 6.8, at 497-500.

industry actors and across different decisionmakers that are unavoidable under a standards-based approach.[508]

Of the approaches discussed in this Article, only the all-or-nothing approaches and the disclosure rule have lower costs of regulation. They, however, either are 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 field against those—end users, low-cost innovators, or start-ups, nonprofits, independent artists, and members of underserved communities—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 chance in the marketplace—that they will be able to reach users and compete with other applications on the merits, without interference from network providers.[509]

In addition to the costs of the nondiscrimination rule itself, there will be costs from 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 that would allow network management as long as it is "reasonable" without further specifying the term,[510] the proposed exception clearly specifies

---

508. *See generally id.* § 6.8, at 500-01.

509. On the importance of certainty for network providers and application developers, see Part II.C.4.a above.

510. 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."

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.[511]

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,[512] 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 nondiscrimination 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.[513]

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 nondiscrimination rule are considerably lower than those of most of the other proposed nondiscrimination 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.

Appropriate Framework for Broadband Access to the Internet over Wireline Facilities, 20 FCC Rcd. 14,986, 14,988 n.15 (2005); Preserving the Open Internet, 24 FCC Rcd. 13,064, 13,113 (proposed Oct. 22, 2009).

    511. These exceptions are afflicted with the same problems as nondiscrimination rules that use ambiguous or undefined terms to describe which discriminatory behavior is banned. *See supra* Parts II.C.3, II.C.4.

    512. *See supra* note 482 and accompanying text.

    513. 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, Background Paper, *supra* note 29, at 10-11. Like all questions related to charging for Quality of Service, this question is outside the scope of this Article.

iii.   *The proposed rule in practice: questions and answers*

In the past four years, I have presented the proposed nondiscrimination rule in many different forums.[514] A number of questions regarding the rule's relationship with Quality of Service have come up again and again, so I answer them here.

1. Does the proposed 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 above.[515]

2. Does the proposed rule require Internet service providers to offer Quality of Service end-to-end—that is, between the original source and ultimate destination of data—across the networks of different providers ("end-to-end, interprovider Quality of Service")?

The rule does not require Internet service providers to offer Quality of Service end-to-end.[516] Network neutrality rules usually have a limited scope, and the obligations imposed by the nondiscrimination 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.[517] 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),"[518] 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."[519]

3. How does the proposed rule constrain the provision of end-to-end, interprovider Quality of Service?

Although network neutrality rules may not cover the entire path of traffic between two endpoints and therefore cannot influence network providers' ac-

---

514. A number of them are listed in the acknowledgements.

515. *See* text accompanying notes 475-82.

516. 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.

517. 47 C.F.R. § 8.11 (2014); *Open Internet Order*, 25 FCC Rcd. 17,905, 17,932-36 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

518. 25 FCC Rcd. at 17,933.

519. *Id.* at 17,933 n.150.

tions along the entire path, network neutrality rules constrain the provision of end-to-end, interprovider 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, interprovider offering.

Although end-to-end, interprovider Quality of Service has not been widely deployed in the public Internet for technical and commercial reasons,[520] any network neutrality nondiscrimination rule should not prevent a migration to end-to-end, interprovider Quality of Service in the future. My proposal reflects this goal. While it constrains the ability of a provider of last-mile Internet access to charge for the provision of Quality of Service, it does not affect payments among interconnecting networks that are not last-mile networks. If a subscriber ("user *A*") requests a certain class of service for particular traffic (whether upstream or downstream), my proposed rules would allow its Internet service provider ("ISP *A*") to charge only its own subscriber, user *A*, for the provision of that class of service on ISP *A*'s network. ISP *A* is not allowed to charge the end user at the other end of the connection (e.g., an application or content provider that is not a customer of ISP *A*'s Internet access service) for the provision of Quality of Service on ISP *A*'s network.[521] To prevent last-mile Internet service providers from exploiting their terminating access monopoly over access to their subscribers, last-mile Internet service providers are not allowed to charge interconnecting networks, content delivery networks, or application providers that would like to interconnect directly with a last-mile Internet service provider for interconnection regarding the transport of upstream and downstream traffic to their subscribers, either, regardless of whether that inter-

---

520. For an analysis of the reasons end-to-end, interprovider Quality of Service has not been widely deployed, see, for example, Grenville J. Armitage, *Revisiting IP QoS: Why Do We Care, What Have We Learned?*, 33 ACM SIGCOMM COMPUTER COMMS. REV. 81, 81-88 (2003); Bell, *supra* note 166; L. Burgstahler et al., *Beyond Technology: The Missing Pieces for QoS Success*, 2003 PROC. ACM SIGCOMM 2003 WORKSHOPS 121; Huston, *supra* note 124.

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.*, VON BORNSTAEDT, *supra* note 171; Verizon Presentation, *supra* note 171, at 41-44; *Verizon Private IP Service Differentiators*, VERIZON (2010), http://www22.verizon.com/wholesale/attachments/solutions/W5036.pdf.

521. *See supra* note 477 and accompanying text.

connection involves the provision of Quality of Service or not. However, my proposed rules would allow, but not require, last-mile Internet service providers to pay interconnecting networks for interconnection (e.g., to buy transit), including for the provision of Quality of Service. Those networks do not have a terminating access monopoly, so payments by last-mile Internet service providers to them do not create the same problems as payments by them to last-mile Internet service providers.[522]

4. If the proposed rule does not require Internet service providers to offer end-to-end, interprovider Quality of Service, and end-to-end, interprovider Quality of Service is not currently widely deployed in the public Internet, what, if anything, 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 United States and Europe, the access networks are the main sources of congestion, while the backbone is not congested.[523] 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 ser-

---

522. As explained in note 480 and accompanying text above, the restrictions on charging for Quality of Service do not flow from the nondiscrimination rule and would have to be encoded separately. Under the network neutrality regime that I propose, the restrictions on charging application providers that are not an Internet service provider's Internet access customers would be captured by a prohibition on all forms of access fees. My network neutrality regime would also include a rule that prohibits last-mile Internet service providers from charging interconnecting networks, content delivery networks, or application providers that interconnect directly with last-mile Internet service providers for the termination and transport of their traffic to and from that Internet service provider's subscribers. On access fees, see note 62 (defining the term) above as well as Box 2 and notes 29-30 and accompanying text above. These questions are outside the scope of this Article.

523. Bauer et al., *supra* note 133, at 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).

There may be other points of congestion. For example, the links that connect interconnecting networks are often congested, too.

vice for upstream 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.[524] 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. 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.[525] 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.[526]

5. Under the proposed 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 megabits of low-delay service over a certain period of time (e.g., per second).[527] 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

---

524. On providing Quality of Service for upstream and downstream traffic, see Question 9 below.

525. Carpenter & Nichols, *supra* note 138, at 1481; Steven Blake et al., Internet Eng'g Task Force, RFC 2475, An Architecture for Differentiated Services (Aug. 1998), https://tools.ietf.org/html/draft-ietf-diffserv-arch-01.

526. Carpenter & Nichols, *supra* note 138, at 1484, 1487-88.

527. 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, *supra* note 15, at 657-59.

on low-delay traffic for that subscriber and re-marks any excess packets according to an agreed upon scheme. For example, the provider may re-mark packets that exceed the maximum limit on low-delay traffic as best-effort service.[528] 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.

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 megabits of low-delay service per second for Internet telephony, but only y megabits 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.[529] 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, for example, 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.

---

528. 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 Presentation, *supra* note 171, at 29. On policing and Quality of Service more generally, see KUROSE & ROSS, *supra* note 15, at 650-52, 657-59, 663; and PETERSON & DAVIE, *supra* note 15, at 550.

529. This practice is called oversubscription. For a short explanation of the practice, see St. Arnaud CRTC Report, *supra* note 357, ¶¶ 8-15.

7. Would the proposed 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 it does not have them, it rejects it. Admission control prevents situations in which all flows are admitted to the network but none receives the amount of resources necessary to perform adequately.[530]

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 the proposed rule, can the provision of Quality of Service ever be based on application-specific criteria?

The proposed nondiscrimination 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.[531]

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.[532] Network providers sell business customers the option to choose the level of

---

530. On admission control in Quality of Service architectures, see KUROSE & ROSS, *supra* note 15, at 665-69; and PETERSON & DAVIE, *supra* note 15, at 539, 542-43.

531. 47 C.F.R. § 8.9(a) (2014) ("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."); *Open Internet Order*, 25 FCC Rcd. 17,905, 17,963-64 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

532. This statement is based on many conversations with networking experts. *See also* Reply Comments of the Center for Democracy & Technology, *supra* note 452, at 20-21; Center for Democracy & Technology Comments, *supra* note 36, at 27; Jordan & Ghosh, *supra* note 366, at 12:21.

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.[533]

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.

For this to work, (1) the user must be able to express his 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.[534] 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.[535]

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—for example, the application, the user's end host, or a home router—could directly

---

533. Verizon Presentation, *supra* note 171, at 25, 29, 56. For a prototype targeting home users, see Yiakoumis et al., *supra* note 434, at 1116-18, which won the Grand Prize at the Imagine App Challenge at the Cable Show 2012. *See* Silbey, *supra* note 478; Spangler, *supra* note 478.

534. 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., *supra* note 525, at 2-5; K. Nichols et al., Internet Eng'g Task Force, RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (Dec. 1998), https://www.ietf.org/rfc/rfc2474.txt.

535. As explained in the main text of the Article, 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., *supra* note 434.

set the Quality of Service-related bits in affected packets in line with the user's choice as expressed through that interface.[536]

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 domain over which the user's Internet service provider offers Quality of Service.[537] 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.[538] The Center for Democracy and Technology described another potential mechanism in its reply comments in the FCC's Open Internet proceeding.[539]

---

536. Applications, end hosts, or home routers could all expose an interface for selecting classes of service, and each could set the appropriate QoS-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 Article.

537. 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 provider 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.

538. *See* R. Hancock et al., Internet Eng'g Task Force, RFC 4080, Next Steps in Signaling (NSIS): Framework (June 2005), https://tools.ietf.org/html/rfc4080; J. Manner et al., Internet Eng'g Task Force, RFC 5974, NSIS Signaling Layer Protocol (NSLP) for Quality-of-Service Signaling (Oct. 2010), https://tools.ietf.org/html/rfc5974.

539. Reply Comments of the Center for Democracy & Technology, *supra* note 452, at 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

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.[540]

In the examples discussed so far, the user makes her 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 interface (e.g., in the form of a website) 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 the proposed nondiscrimination rule, but I am not convinced that network neutrality rules should allow Quality of Service. Can I adopt the proposed nondiscrimination rule but ban Quality of Service? And if I do so, would the proposed nondiscrimination rule still be useful?

Nondiscrimination rules apply to any form of differential treatment, not just to the differential conduct necessary to provide Quality of Service.[541] Nondiscrimination rules govern any differential handling of packets within the network, including, for example, the allocation of bandwidth among users during times of congestion or other differential treatment of data packets to manage congestion. Nondiscrimination rules also regulate differential pricing practices directed towards subscribers[542]: they affect, for example, whether network providers can count only some types of traffic but not others towards users' monthly usage caps, what factors network providers can use to price discrimi-

---

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.").

540. Yiakoumis et al., *supra* note 434, at 1117 ("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."); E-mail from Yiannis Yiakoumis to author (Mar. 20, 2013) (on file with author) ("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).").

541. For a more detailed description of the scope of nondiscrimination rules, see Part II.A above.

542. The nondiscrimination rules discussed in this Article do not address a network provider's pricing practices with respect to application providers that are not its Internet service customers. Whether an Internet service provider can charge application providers who are not its subscribers an "access fee" for access or for prioritized or otherwise enhanced access to its subscribers is governed by a network neutrality regime's rules about access fees. For a definition of the term "access fees," see note 62 above. A discussion of access fees is outside the scope of this Article. For a short overview of the debate, see Box 2 and notes 29-30 above.

nate, and 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 nondiscrimination rule has important implications for many questions other than whether and under which conditions to allow Quality of Service.

The nondiscrimination rule proposed in the text allows certain forms of user-controlled Quality of Service. Regulators who like the rule but disagree with the Article's arguments regarding Quality of Service could adopt the nondiscrimination rule proposed here and couple it with an explicit ban on Quality of Service. In this case, the nondiscrimination rule would apply to any form of differential treatment except for Quality of Service. Thus, the choice of nondiscrimination rule can be separated from the decision whether to allow Quality of Service if that is desired.

III.   THE OPEN INTERNET ORDER'S NONDISCRIMINATION RULE

In October 2009, the FCC 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 nondiscrimination rule that would have required network providers to treat every packet the same, subject to reasonable network management.[543]

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.[544] All of the proposals for nondiscrimination rules discussed in this Article were supported by some commenters.[545]

In December 2010, FCC Chairman 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 Representative Waxman, Chairman of the House Committee on Energy and Commerce, and Representative Boucher, 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.[546] Not all participants in the negotiations supported

---

543. *See supra* note 107.

544. *Open Internet Order*, 25 FCC Rcd. 17,905, 17,906 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

545. See the sources cited in support of the different proposals above.

546. *See* Kim Hart, *Sources: FCC Chief to Move on Net Neutrality*, POLITICO (Nov. 18, 2010, 4:54 PM EST), http://www.politico.com/news/stories/1110/45371.html. 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. AT&T, Notice of Ex Parte Conversations, Preserving the Open Internet, GN Docket No. 09-191, Framework for Broadband Internet Service, GN Docket No. 10-127 (Nov. 22, 2010), *available at* http://apps.fcc.gov/ecfs/document/view?id=7020921306; Todd Shields, *AT&T Gains FCC's Ear as Regulators Near Decision on Net Neutrality Rules*, BLOOMBERG (Nov. 30,

the final proposal.[547] The proposal was never introduced. Representative Waxman and Representative Boucher abandoned the effort when the Republican members of the House Committee on Energy and Commerce refused to support the bill.[548]

The rules proposed by the Chairman included a nondiscrimination rule that applied to fixed, but not mobile, Internet access service. Like the nondiscrimination rule in the Waxman proposal, the rule banned "unreasonable" discrimination without specifying how to interpret the term and left it to later case-by-case adjudication to decide whether specific discriminatory conduct meets this criterion.[549]

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,[550] that rule does not adequately protect the values that network neutrality regulations are designed to protect.[551] Thus, from the perspective of network neutrality proponents, this proposal was highly unsatisfactory.[552]

2010), http://www.bloomberg.com/news/2010-11-30/at-t-gains-fcc-s-ear-as-regulators-near -decision-on-net-neutrality-rules.html.

The Open Internet Rules regarding blocking, nondiscrimination, 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 Open Internet Order*, 25 FCC Rcd. at 17,992-93, *with* Draft Bill, *supra* note 230. For more on the Waxman proposal, see note 230 above.

547. In particular, Free Press and the Open Internet Coalition did not support the compromise proposal. *See* Eggerton, *supra* note 230; Jerome, *OIC*, *supra* note 230. 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* Public Knowledge & Media Access Project, Notice of Oral Ex Parte Presentation, Broadband Industry Practices, WC Docket No. 07-52, Preserving the Open Internet, GN Docket No. 09-191, Framework for Broadband Internet Service, GN Docket No. 10-127 (Dec. 10, 2010), *available at* http://apps.fcc.gov/ecfs/document/view?id =7020923323 (disclosing a meeting in which representatives of the 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).

548. Hart, *supra* note 230; *see also* Press Release, Henry A. Waxman, *supra* note 230.

549. The draft rules were not released publicly, but they were 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, *supra* note 229. For the nondiscrimination rule in Waxman's proposal, see Draft Bill, *supra* note 230, § 2, at 2 (proposing to add § 12(a)(1)(B)). The nondiscrimination 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 id.* § 2, at 4 (proposing to add § 12(c)).

550. *See supra* notes 280-88 and accompanying text.

551. *See supra* Parts II.C.3, II.C.4.

552. Network neutrality proponents also opposed other aspects of the proposal. For an overview, see van Schewick, *supra* note 406; and van Schewick, *supra* note 263.

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.[553] 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.[554]

This put the Chairman in a difficult position. During his presidential campaign, President Obama had promised to enact network neutrality rules.[555] Attempts to enact network neutrality rules in Congress had failed.[556] 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.[557]

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.[558] The Chairman and the two Democratic Commissioners negotiated

---

553. For a summary of their criticisms, see *Open Internet Order*, 25 FCC Rcd. 17,905, 18,046 (2010) (Copps, Comm'r, concurring); *id.* at 18,082-83 (Clyburn, Comm'r, approving in part and concurring part); and David Hatch, *FCC's Copps Isn't Committed on Net Neutrality*, Nat'l J. Daily (Dec. 9, 2010, 3:30 PM), http://www.nationaljournal.com/member /daily/fcc-s-copps-isn-t-committed-on-net-neutrality-20101209.

554. Hatch, *supra* note 553.

555. *Barack Obama*, *supra* note 282.

556. The last such attempt had failed in September 2010, when the Republican leadership in Congress refused to support the compromise proposal negotiated by Representative Waxman and Representative Boucher. Hart, *supra* note 230; *see also* Press Release, Henry A. Waxman, *supra* note 230.

557. *See Despite Seat at Table, AT&T Outed Delicate Net-Neutrality Talks*, Hill (Nov. 25, 2010, 9:52 PM ET), http://thehill.com/policy/technology/130737-despite-seat-at-the -table-atat-outed-delicate-net-neutrality-talks-sparking-gop-backlash [hereinafter *Seat at Table*]; Hart, *supra* note 546.

558. At the time, AT&T was the second-largest provider of wireline Internet access in the United States, the largest DSL provider in the United States, and the second-largest wireless provider in the United States. *See* Press Release, Leichtman Research Grp., Over 800,000 Add Broadband in the Third Quarter of 2010 (Nov. 15, 2010), http://www .leichtmanresearch.com/press/111510release.html; *see also Market Share of Mobile Wireless Telecommunication Subscribers by Selected Service Providers in the United States at the End of 2008 and 2009*, Statista, http://www.statista.com/statistics/184755/market-share-of -mobile-wireless-subscribers-by-providers-in-the-us (last visited Jan. 7, 2015); *Wireless Market Share 2011-2014*, *supra* note 288. AT&T, which has consistently outspent every other corporation in donations to congressional campaigns in the past, has considerable influence in Washington. *See, e.g.*, AT&T, Notice of Ex Parte Conversations, *supra* note 546 (disclosing two conversations between Jim Cicconi, Senior Executive Vice President of External and Legislative Affairs at AT&T and head of AT&T's public policy organization, and Edward Lazarus, the Chairman's Chief of Staff, in which "Mr. Cicconi discussed the merits

over improvements to the Order up to the day before the Commission's open meeting on December 21.[559] In the end, they agreed on the following compromise[560]: They did not change the text of the nondiscrimination 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.[561] The Rules were appealed by Verizon, Free Press, and others and, except for the disclosure rule, vacated by the Court of Appeals for the D.C. Circuit in January 2014.[562]

In the final Open Internet Report and Order, the FCC adopted a nondiscrimination rule that banned providers of fixed broadband Internet access service from "unreasonably discriminat[ing] in transmitting lawful network traffic

---

of the proposed Waxman legislation and why it should be a model for a substantive resolution of the issues raised in these proceedings"); *Seat at Table*, *supra* note 557 ("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.'"); Cecilia Kang, *AT&T, Carriers Fund Democratic Reps Against Net Neutrality*, WASH. POST POST TECH (Oct. 19, 2009, 9:00 AM ET), http://voices.washingtonpost.com/posttech/2009/10/all_but_two_of_the.html (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 52 out of the 72); Matthew Lasar, *AT&T and Astroturf: Is "Following the Money" Enough?*, ARS TECHNICA (Oct. 26, 2009, 7:30 PM PDT), http://arstechnica .com/tech-policy/2009/10/the-anti-net-neutrality-movement-is-it-just-about-att-money; Matthew Lasar, *Hey, Capitol Hill: Who's Your Daddy? AT&T*, ARS TECHNICA (Sept. 14, 2010, 4:30 AM PDT), http://arstechnica.com/tech-policy/2010/09/hey-capitol-hill-whose-your -daddy-att; Lindsay Renick Mayer, *Big Donors Ramp Up to Fight FCC Net Neutrality Decision*, OPENSECRETS.ORG (Oct. 28, 2009, 7:26 PM), http://www.opensecrets.org/news /2009/10/the-federal-communications-com.html; Editorial, *The Price of Broadband Politics*, N.Y. TIMES (June 29, 2010), http://www.nytimes.com/2010/06/30/opinion/30wed2.html; Shields, *supra* note 546 (discussing a large number of conversations between Cicconi and Lazarus 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 to the rules and protect the Commission against Republican criticism).

559. Grant Gross, *Net Neutrality Plan Has the Votes at FCC*, COMPUTERWORLD (Dec. 20, 2010, 5:14 PM PT), http://www.computerworld.com/s/article/9201818/Net_neutrality _plan_has_the_votes_at_FCC.html; Arik Hesseldahl, *Breaking: FCC's Copps Voting "Yes" on Net Neutrality Plan*, ALL THINGS DIGITAL (Dec. 20, 2010, 1:45 PM PT), http://allthingsd .com/20101220/breaking-fcc-commissioner-copps-says-hell-vote-yes.

560. This Article focuses on nondiscrimination rules. For an overview of the other parts of the compromise, see Barbara van Schewick, *The FCC's Open Internet Rules—Stronger than You Think*, INTERNET ARCHITECTURE & INNOVATION (Dec. 27, 2010), http:// netarchitecture.org/2010/12/the-fcc's-open-internet-rules---stronger-than-you-think.

561. Preserving the Open Internet, 76 Fed. Reg. 59,192, 59,223 (Sept. 23, 2011) (codified at 47 C.F.R. pts. 0, 8).

562. Verizon v. FCC, 740 F.3d 623, 659 (D.C. Cir. 2014); Notice of Multicircuit Petitions for Review, Preserving the Open Internet, GN Docket No. 09-191, Broadband Industry Practices, WC Docket No. 07-52 (J.P.M.L. Oct. 5, 2011), *available at* http://hraunfoss.fcc .gov/edocs_public/attachmatch/DOC-310240A1.pdf.

over a consumer's broadband Internet access service" and stipulated that "[r]easonable network management shall not constitute unreasonable discrimination."[563] The FCC was to determine whether certain discriminatory conduct qualified as unreasonable in case-by-case adjudications.

The text of the Order specifies the factors that the FCC would have used to determine whether certain discriminatory conduct constitutes unreasonable discrimination[564]: 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."[565]

Use-agnostic discrimination (or "application-agnostic" discrimination), the Order explains, is "[d]ifferential treatment that does not discriminate among specific uses of the network or classes of uses."[566] 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.[567] This is the same substantive standard as the one used by the nondiscrimination rule proposed by this Article. (See Table 2: The Open Internet Order's Nondiscrimination Rule and the Rule Proposed by This Article below.) As explained above, allowing use-agnostic discrimination but banning discrimination among uses or classes of uses preserves the application-agnosticism of the network.[568]

TABLE 2
The Open Internet Order's Nondiscrimination Rule and
the Rule Proposed by This Article

	Application-Agnostic Discrimination	Application-Specific Discrimination
FCC's Nondiscrimination Rule	Likely to be reasonable	Likely to be unreasonable
Nondiscrimination Rule Proposed by This Article	Allowed	Banned

---

563. *Open Internet Order*, 25 FCC Rcd. 17,905, 17,992 (2010) (report and order), *vacated in part*, *Verizon*, 740 F.3d 623.

564. *Id.* at 17,944-46. The following paragraphs are based in part on van Schewick, *supra* note 560.

565. 25 FCC Rcd. at 17,946.

566. *Id.* at 17,945.

567. *Id.* at 17,945-46.

568. *See supra* Box 19; *supra* note 459 and accompanying text.

Thus, in evaluating whether discriminatory conduct is reasonable, the FCC was to 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.[569] 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."[570] 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 (the third of the four factors that have fostered application innovation in the past) as explicit purposes of the Rules.[571]

As this Article has shown, using these factors—application-agnosticism, user choice, and innovation without permission[572]—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 Article would be evaluated under the "no unreasonable discrimination" standard, the analysis in this Article suggests which results the FCC would reach if it takes these factors seriously. (See Table 3: Evaluating Different Forms of Quality of Service Under the Open Internet Order's Nondiscrimination 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

---

569. On these factors (the other two being innovation without permission and low cost of application innovation), see Boxes 3 and 4 and notes 59-62 and accompanying text above.

570. 25 FCC Rcd. at 18,046 (Copps, Comm'r, concurring); *see also id.* at 17,944-45 (report and order) ("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.'" (footnote omitted)); *id.* at 17,946 ("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.").

571. 47 C.F.R. § 8.1 (2014) ("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 25 FCC Rcd. at 17,907-10.

572. 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 nondiscrimination rule.

provider is allowed to charge only its own Internet service customers for the use of the different classes of service[573] is likely to be reasonable.[574] By contrast, offering Quality of Service exclusively to one or more applications within a class of "like" applications is unlikely to be reasonable.[575] 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 discriminate among classes of applications that are alike and does not discriminate among applications within a class of like applications).[576]

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.[577] The Order explicitly rejects the view that the nondiscrimination rule should only prohibit discrimination that is "anticompetitive."[578] As explained above, such a rule (or an interpreta-

---

573. 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 nondiscrimination 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 about these arrangements, 25 FCC Rcd. at 17,919-25, 17,947-48, unequivocally rejects the main arguments in favor of them, *id.* at 17,921-22, 17,928-29, and concludes that "as a general matter," arrangements of this kind are "unlikely" to be considered reasonable under the nondiscrimination rule, *id.* at 17,947-48. Thus, under the Open Internet Rules, network providers (most likely) would have been allowed to charge only their own Internet access service customers for any differential treatment allowed by the nondiscrimination rule. For a more detailed analysis, see van Schewick, *supra* note 560.

574. *See supra* notes 476-86 and accompanying text; *infra* Table 3.

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.

25 FCC Rcd. at 17,944-45.

575. *See supra* notes 372-81 and accompanying text; *infra* Table 3.

576. *See supra* notes 383-442 and accompanying text; *infra* Table 3.

577. See the sources cited throughout Part II.C.1 above.

578. 25 FCC Rcd. at 17,949-50 ("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 rea-

tion 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 threatens the values network neutrality rules are designed to protect.[579]

Finally, according to the Order, the same principles that guide the Commission's interpretation of the nondiscrimination rule were also to guide the Commission's evaluation of network management practices under the Open Internet Rules' exception for reasonable network management.[580] The exception applied to the no-blocking rules for fixed and mobile Internet access and to the nondiscrimination 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.[581] 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.[582] In line with these considerations, the Order made clear that network management would be evaluated by the same principles that guide the interpretation of the nondiscrimination rule.

Overall, the nondiscrimination rule adopted by the Commission (as clarified by the text of the Order) constituted a considerable improvement over the same rule without clarifications. The general theoretical framework underlying the Order as well as the specific factors that would have been used to interpret the nondiscrimination rule and the reasonable network management exception are in line with the broader theoretical framework on which calls for network neutrality regulation are based.[583] In contrast to the standards used by other case-by-case approaches (e.g., by an antitrust framework or the Verizon-Google legislative framework), the factors the FCC would have used to evaluate differential treatment do not automatically exclude instances of discrimination that threaten the values network neutrality rules are intended to protect.[584] 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 that may end up enforcing network neutrality rules.

Still, compared with the bright-line nondiscrimination rule supported by this Article, considerable uncertainty remains. For example, it is not clear how

---

sons that may be independent of providers' competitive interests or that may not immediately or demonstrably cause substantial consumer harm." (footnotes omitted)).

579. *See supra* Part II.C.1.

580. 25 FCC Rcd. at 17,954.

581. AT&T Open Internet Comments, *supra* note 355, at 183-87; Cox Comments, *supra* note 366, at 30-33.

582. *See, e.g.*, van Schewick, Official Testimony, *supra* note 68, at 4-8, *see also supra* notes 213-16, 251-58, 268-75.

583. *See supra* Part I.

584. *See supra* Parts II.C.1, II.C.2.

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.[585]

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 in evaluating . . . reasonableness."[586] This suggests that, under such a rule, 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 are application-specific and threaten user choice.

---

585. For a discussion of these costs, see Part II.C.4 above.

586. *Open Internet Order*, 25 FCC Rcd. 17,905, 17,944-46 (2010) (report and order), *vacated in part*, Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014).

TABLE 3
Evaluating Different Forms of Quality of Service Under
the Open Internet Order's Nondiscrimination Rule

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 available equally to all applications and classes of applications; (2) the user is able to choose whether, when, and for which applications 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 nondiscrimination rule	No	No	Yes

Note: the "Types of QoS" header label appears in the top-left diagonal cell.

* According to 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 would be allowed to charge only their own Internet access service customers for any differential treatment allowed by the nondiscrimination rule. *See supra* note 573.

TABLE 4
The Impact of Different Nondiscrimination Rules on Quality of Service

Rules \ Forms of QoS	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 available equally to all applications and classes of applications; AND (2) user is able to choose whether, when, and for which applications to use which class of service*
Ban application-specific discrimination, allow application-agnostic discrimination	banned	banned	allowed, if (1) different classes of service are available equally to all applications and classes of applications; AND (2) user is able to choose whether, when, and for which applications to use which class of service*
FCC's nondiscrimination rule	likely banned	likely banned	likely allowed, if (1) different classes of service are available equally to all applications and classes of applications; (2) user is able to choose whether, when, and for which applications 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 policymakers adopt a nondiscrimination 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 nondiscrimination rule. *See supra* note 80.

** According to 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 would be allowed to charge only their own Internet access service customers for any differential treatment allowed by the nondiscrimination rule. *See supra* note 573.

CONCLUSION: NETWORK NEUTRALITY AND QUALITY OF SERVICE

The network neutrality debate is often framed as a debate for or against Quality of Service. As this Article 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.[587] 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 nondiscrimination rule affects network providers' ability to offer Quality of Service. To address this problem, this Article 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 Nondiscrimination Rules on Quality of Service above.

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, this Article offers interesting new insights.

Most proponents of network neutrality 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 Article shares that view.[588] 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."[589] 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 differently from Skype, or Gmail differently from Hotmail.[590] In the United States, 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 al-

---

587. On these proposals, see Parts II.B.1 and II.B.2 above.

588. On this form of Quality of Service, see Part II.D.2.a.i above.

589. On this form of Quality of Service, see Part II.D.2.a.ii above.

590. Internet telephony is sensitive to delay, but e-mail is not, so the two classes of applications are not alike. *See supra* note 15.

lowed. As this Article 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.

By contrast, Quality of Service architectures *in which network providers make different types of service available equally to all applications and classes of applications and in which users choose whether, when, and for which appli-*

cations to use which type of service do not raise similar concerns.[591] *First*, they preserve the *application-agnosticism* of the network: The provision of Quality of Service is dependent not on which applications users are using but on the QoS-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 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 nondiscrimination rule proposed by this Article and the Open Internet Order's vacated nondiscrimination rule for fixed broadband Internet access, these are the only forms of Quality of Service that network providers would be able to offer.

If policymakers adopt a nondiscrimination 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, policymakers 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 nondiscriminatory basis; or (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 that should be considered and evaluated separately. I take

---

591. On this type of Quality of Service, see Part II.D.2.b.i.A above.

up restrictions on charging elsewhere.[592] There, I argue that network providers should only be allowed to charge their own Internet service customers for any differential treatment allowed by the nondiscrimination rule.[593]

Finally, a network provider that 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.[594]

Opponents of network neutrality regulation have created the impression that policymakers 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 this Article 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 infrastructure to evolve. Thus, regulators can have their cake and eat it, too.

---

592. *See* sources cited *supra* note 29.

593. This restriction also applies to interconnection agreements, as discussed in response to Question 3 in Part II.D.2.b.iii. *See also supra* note 480.

594. For a discussion of this requirement, *see* van Schewick, Background Paper, *supra* note 29, at 10-11. The European Union has adopted a similar rule following its review of the regulatory framework for telecommunications services. *See* Universal Service Directive, *supra* note 291, art. 22(3); *Impact Assessment*, *supra* note 293, at 92, 95-97, 101.