Barbara van Schewick
Associate Professor of Law and (by Courtesy) Electrical Engineering
Helen L. Crocker Faculty Scholar
Director, Center for Internet and Society
Stanford Law School

July 31, 2012

Comments on BEREC's Public Consultation on Network Neutrality, BoR (12) 32 - Guidelines for Quality of Service in the scope of Net Neutrality, BoR (12) 31 - Differentiation practices and related competition issues in the scope of Net Neutrality

I welcome the opportunity to comment on BEREC's Public Consultation on Network Neutrality. I submit these comments as a professor of law and, by courtesy, electrical engineering at Stanford University whose research focuses on Internet architecture, innovation and regulation. My book "Internet Architecture and Innovation," which was published by MIT Press in 2010, is considered the seminal work on the science, economics and politics of network neutrality. My papers on network neutrality have influenced discussions on network neutrality all over the world.[1] I have testified on matters of Internet architecture, innovation and regulation before the US Federal Communications Commission.[2] I have not been retained or paid by any of the parties to this consultation.

My recent paper "Network Neutrality and Quality of Service: What a Non-Discrimination Rule Should Look Like,"[3] which I attach, directly addresses many of the issues raised by this consultation. The paper explores the relationship between network neutrality, non-discrimination rules and Quality of Service in detail. In particular, the paper:

- offers the first in-depth analysis of the relationship between network neutrality and Quality of Service;
- proposes a framework that policy makers can use to evaluate alternative proposals for network neutrality rules and to assess specific forms of discriminatory conduct;
- evaluates existing proposals for non-discrimination rules and proposes a non-discrimination rule that policy makers should adopt around the world – a rule that the Federal Communications Commission adopted at least in part; and
- provides the first detailed analysis of the Federal Communications Commissions' non-discrimination rule and of its implications for network providers' ability to manage their networks and offer Quality of Service.

[1] van Schewick (2007); Frischmann & van Schewick (2007).
[2] See, e.g.,van Schewick (2008); van Schewick (2010c); van Schewick (2010b).
[3] van Schewick (2012).

In the process, the paper explains

- why non-discrimination rules that are *based on an antitrust framework* or *ban only behavior that is anticompetitive* do not sufficiently protect users and innovators against all forms of discrimination that network neutrality proponents are concerned about; and
- why non-discrimination rules that *ban discriminatory conduct that is not disclosed* do not adequately protect the values that network neutrality regulation is designed to protect, even in markets where consumers have a choice of more than one Internet service provider. This insight is particularly relevant for the debate over wireless network neutrality in the US and for the network neutrality debate in Europe, Canada or Australia.

The following comments highlight some of the key insights of the paper that may be useful for BEREC in considering this topic.[4] I also add some additional comments that draw on my earlier work.

1. BEREC's thinking about network neutrality should be guided by the underlying goals of network neutrality regulation.

BEREC's thinking about network neutrality is, of course, driven by the existing regulatory framework in the EU. This framework should, however, be interpreted and applied in light of the underlying goals of network neutrality regulation. As I explain in the attached paper, 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 non-economic ways. There is, however, a lot of uncertainty on how to get from a high-level commitment to network neutrality to a specific set of rules. The paper proposes a framework that policy makers and others can use to interpret existing network neutrality rules in light of the goals of network neutrality regulation or to decide whether specific discriminatory conduct should or should not be allowed under a network neutrality regime. According to that framework, network neutrality rules should meet the following criteria:

- Network neutrality rules should preserve the factors that have allowed the Internet to serve as a platform for application innovation and economic growth and as a platform for free speech and decentralized economic, social, cultural and political interaction in the past. As I explained elsewhere, these factors are:[5]

[4] Throughout these comments, I draw heavily on the attached paper.

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

- *Innovator choice:* 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 (this factor has also been called "*innovation without permission"*). Adding additional decision-makers who need to endorse the idea or take action before an idea can be realized reduces the chances that innovative ideas can be realized.[6]
- *User choice:* Users independently choose which applications they want to use, without interference from network providers. Letting users, not network providers choose which applications will be successful is an important part of the mechanism that produces innovation under uncertainty.[7] 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.[8] (See *Box 4: The Importance of User Choice* on p. 9 of the attached paper.)
- *Application-Blindness:* The network is application-blind. An application-blind network is unable to distinguish among the applications on the network, and, as a result, is unable to make distinctions among data packets based on this information.[9] The application-blindness of the network ensures that network providers cannot interfere with innovators' and users' choices, that they cannot distort competition among applications (or classes of applications) or reduce application developers' profits through access fees[10] (we may call this "*innovation without fear"*).
- *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.[11] If there is uncertainty (e.g., about

---

(2001); Lessig (2002); Wu (2003); Wu & Lessig (2003); van Schewick (2004); Wu (2004); Cerf (2006); Lessig (2006); Lessig (2008).

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

[7] See van Schewick (2010a), pp. 349-351; van Schewick (2010c), p. 6. See also footnote 52 in the attached paper.

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

[9] The original Internet was application-blind, which was a consequence of its architecture, in particular of the broad version of the end-to-end arguments and of the layering principle. See footnote 2 in the attached paper and van Schewick (2010a), pp. 72-75, 217-218; van Schewick (2004). See also, e.g., Cerf (2006), pp. 1-4, 7; Reed (2010). For a short summary of the importance of application-blindness, see van Schewick (2010c), pp. 3-4. For a detailed analysis, see van Schewick (2010a), pp. 215-281, 286-295, 349-353, 355-365. See also Benkler (2000), pp. 565-568; Balkin (2009); van Schewick (2010a), pp. 359-365 (all focusing on the social, cultural and political implications).

[10] 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. On access fees, see, e.g., van Schewick (2010b). See also footnotes 15 to 20 in the attached paper and accompanying text.

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

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.[12] In the current Internet, there is uncertainty and user needs are heterogeneous, so the conditions under which innovator diversity increases the amount and quality of innovation are met.[13]

- Network neutrality rules should not constrain the evolution of the network more than is necessary to reach the goals of network neutrality regulation.
- Network neutrality rules should make it easy to determine which behavior is and is not allowed to provide much-needed certainty for industry participants. For application providers, uncertainty over the level of protection provided by the rule reduces their incentives to innovate and their ability to get funding.[14] For network providers, uncertainty over the legality of network management practices or of different forms of Quality of Service may make it more difficult to manage their network and may limit the evolution of the network infrastructure. Uncertainty over the regulatory regime may also reduce network providers' incentives to invest more generally.[15]
- Network neutrality rules should keep the costs of regulation low.

2. BEREC should aim to specify applicable rules and guidelines in advance to provide much-needed certainty to industry participants.

BEREC's proposed "Guidelines for Quality of Service in the scope of Network Neutrality" currently leave all decisions over the legality of specific discriminatory conduct to future adjudications. As the paper shows, such an approach creates considerable social costs.[16] Such an approach fails to provide much-needed certainty to industry participants. Network providers still will not know which forms of network management are acceptable, which constrains the evolution of the network more than necessary. Application developers will not know in advance against which discriminatory conduct they are protected. This decision will only be made after they have been discriminated against and gone through a long and expensive process. The resulting uncertainty reduces their incentives to innovate and their ability to get funding. Moreover, an approach that decides ex post whether a specific form of discriminatory conduct is allowed creates high costs of regulation and tilts the playing field against those – end users, low-cost application developers and start-ups – who do not have the resources to engage in extended fights over the legality of specific discriminations in the future. Finally, as the paper shows, deciding the legality of specific

[12] For a short version of the argument, see van Schewick (2010c), pp. 5-6 and van Schewick (2010b), pp. 4-5. For the detailed version, van Schewick (2010a), pp. 298-349.
[13] See van Schewick (2010a), p. 356.
[14] See footnotes 174 to 176 in the attached paper and accompanying text.
[15] See footnotes 171 to 173 in the attached paper and accompanying text.
[16] See Section "Problems with Case-by-Case Adjudication" in the attached paper.

discriminatory conduct in individual adjudications is unlikely to lead to decisions that adequately protect the values network neutrality rules are intended to protect.

Given the structure of the regulatory framework, BEREC may not be able to remove all uncertainty. It should, however, aim to provide as much certainty in advance as possible by, for example, clearly specifying the criteria that National Regulatory Agencies will use to evaluate specific forms of differential treatment in the future.

3. BEREC should clarify that it will evaluate all forms of differential treatment based on whether the differential treatment is application-specific or application-agnostic. Application-specific discrimination should be banned, application-agnostic discrimination should be allowed.[17] 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.[18] This general rule should be coupled with an exception for reasonable network management, which would allow narrowly tailored deviations from the non-discrimination rule if a network management problem cannot be addressed in application-agnostic ways.

This approach reinforces key architectural principles on which the Internet was based without locking in the original architecture of the Internet itself. It 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 providing 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.

BEREC's "Guidelines for Quality of Service in the scope of Network Neutrality" already indicate that application-specific traffic management will be considered problematic if application-agnostic ways of solving the specific traffic management problem are available, and that any traffic management measures must be efficient (i.e. finely tuned so as to achieve exactly the pursued

[17] In the Federal Communications Commission's Open Internet Proceeding, this proposal was supported by, e.g., networking experts (e.g., Reed (2010); NYSERNet (2010)); venture capitalists (e.g., Burnham (2010); Wilson (2010)), entrepreneurs (e.g., Borthwick (2010); Srinivasan & Gupta (2010)) and non-profit organizations (e.g., The Council of Scientific Society Presidents (2010); North American Benthological Society (2010); Botanical Society of America (2010)).

[18] Relevant characteristics of an application include what the attached paper calls "application" (i.e. 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. SIP vs. Skype's proprietary protocol, or TCP vs. UDP), or the application's technical requirements (e.g., latency-sensitive vs. non-latency-sensitive applications). Since the term "applications" stands for applications, content, services or uses, the ban on application-specific discrimination applies equally to discrimination based on criteria that depend on characteristics of content or characteristics of a service or use. Thus, discrimination against certain content based on, e.g., publisher, author, content type, subject matter, or viewpoint would all be prohibited by this approach.

objective) and proportionate (i.e. leading to as few side effects as possible.)[19] In this respect, the approach proposed here would lead to the same result.

As I explain in the paper,[20] the approach proposed here should apply to all forms of differential treatment, regardless of the purpose (i.e., to all forms of network management (i.e., not just to congestion management), to the offering of traffic classes, or to other forms of differentiation). As the paper shows, any measure that singles out an application or class of applications for differential treatment tilts the playing field against some applications or classes of applications and interferes with users' decisions about how to use the network, creating significant social costs. At the same time, network providers can usually realize their legitimate goals using application-agnostic means that are not similarly harmful to application innovation, user choice, or the Internet's ability to reach its social, cultural or political potential. Based on these insights, the approach takes away all the tools that would allow network providers to deliberately or inadvertently interfere with competition and user choice – those involving application-specific discrimination –, while leaving the tools that cannot distort competition or violate user choice – those involving application-agnostic discrimination.

By legitimizing a broad range of discriminatory conduct (that is, all conduct that is application-agnostic), the approach gives network providers great flexibility to realize legitimate goals such as congestion management, price discrimination or product differentiation, albeit through means that do not interfere with the values that network neutrality rules are designed to protect. For example, during times of congestion, a network provider could give one person a larger share of the available bandwidth than another, for example because this person pays more for Internet access or has used the Internet less over a certain period of time.[21] But it could not throttle the bandwidth available to a specific online video application (e.g., BitTorrent or YouTube) or to online video in general.[22]

Application-agnostic network management coupled with user-controlled prioritization[23] gives network providers the tools they need to maintain the quality of the Internet experience for all users, even during times of congestion, while preserving the application-blindness of the network and the principle of user choice to the extent possible. Network providers would be able to prevent aggressive users from overwhelming the network and enforce fairness among users by allocating bandwidth among users in application-agnostic ways. But how users use the bandwidth available to them, and whether they would like to give some of their applications priority over others, would be choices left to the users. At the same time, the reasonable-network-management exception provides

[19] BEREC Guidelines, p. 52.
[20] See Section "Ban Application-Specific Discrimination, Allow Application-Agnostic Discrimination" in the attached paper.
[21] That would be application-agnostic discrimination.
[22] That would be discrimination based on application or class of application.
[23] To the extent that applications benefit from relative prioritization at times of congestion, network providers could allow users to choose which applications to prioritize within the user's bandwidth envelope during times of congestion. As long as the ability to prioritize 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 is left to the user, this form of network management would be consistent with the non-discrimination rule proposed above.

a safety valve that allows network providers to react in more application-specific ways if a problem cannot be solved through application-agnostic means.

The proposed 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 (1) the different classes of service are offered equally to all applications and classes of applications; (2) the user is able to choose whether and when to use which class of service; and (3) the network provider is allowed to charge only its own Internet service customers for the use of the different classes of service. [24,25,26]

Finally, by clearly specifying in advance which behavior is and is not allowed, the rule provides certainty to all market participants. Network providers would know how they can manage their networks, and application developers (and their investors) could be sure that they will not be discriminated against.

4. Any form of application-specific differential treatment[27] should trigger a need for regulatory intervention, regardless of the market share of the Internet service provider and of the prevalence of the practice at the market level. BEREC's current proposal overestimates the ability of users to discipline providers and would unduly restrict the availability of unrestricted Internet offerings to the detriment of society.

In the draft Guidelines, BEREC suggests assessing practices that have been determined to be problematic (e.g., because they are application-specific) at the market level to determine whether there is a need for regulatory intervention. This proposal rests on two assumptions that are problematic:

First, BEREC's proposal is based on the assumption that if a network provider unduly discriminates against an application that users would like to use, users can switch to another network provider who does not discriminate against the affected application. The threat of switching, proponents of this approach assume, will discipline providers.

This assumption fails 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.[28] Rules that require network providers to disclose whether and how they interfere with applications and content on their networks reduce the

[24] See also the discussion on of user-controlled Quality of Service on pp. 11-11 of these comments.
[25] This restriction would not constrain interconnection agreements in any way. Thus, payments among interconnecting networks would remain possible.
[26] While the first two conditions directly flow from the proposed non-discrimination rule, the third condition is based on additional considerations and would need to be encoded separately.
[27] This would, of course, be subject to the reasonable network management exception outlined above.
[28] On the arguments in this section, see Section "Ban Discrimination That Is Not Disclosed" in the attached paper.

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 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.

Due to all of these problems, less users would switch providers in response to actual discrimination or exclusion (and, consequently, providers will be less deterred by the treat of switching) than would be necessary for switching to have the desired disciplinary effect. Thus, even if there is competition in the market for Internet access services, disclosure cannot replace substantive regulation as a tool to discipline providers.

While the draft guidelines note some of the reasons that may make the threat of switching less effective in disciplining providers, they do not draw the necessary conclusion – that taken together, the various factors discussed above provide network providers with a substantial degree of market power over their customers that enables them to restrict applications and content on their network without losing too many Internet service customers.

Second, the draft guidelines seem to assume that the existence of restricted offerings is less problematic if there are unrestricted offerings available that users can switch to. Such an interpretation fails to protect all values that network neutrality rules are designed to protect. First, due to the factors just described, the availability of unrestricted offerings is not enough to discipline providers.

Moreover, such an interpretation harms users: as described in detail in the attached paper, even users who would like to use the restricted application and would use it if their Internet service provider was not restricting it, may not switch to the unrestricted Internet service offering, making them worse off than they would be if their Internet service provider was banned from restricting the application.[29]

BEREC's interpretation also harms application innovation. User choice is 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. If restricted and unrestricted offerings co-exist on the Internet, only some users – those who use unrestricted offerings – participate in this mechanism. By contrast, in restricted offerings, network providers, not users, pick winners and losers on the Internet. Whom they pick may be driven by a number of motivations that are not necessarily identical with what users would prefer, leading to applications that their users would not have chosen and forcing their users to engage in an Internet

[29] See Section "Ban Discrimination That Is Not Disclosed" in the attached paper.

usage that does not create the value it could. The threat of discrimination also reduces application developers' incentives to innovate.

Finally, a co-existence of restricted and unrestricted offerings would reduce the size of the potential market that application developers, content providers or service providers have access to. In an unrestricted Internet, innovators have immediate access to everybody connected to the Internet. This has allowed innovators to offer applications, content or services for niche markets, since they can reach niche audiences even if they are spread over many providers.[30] Even for providers targeting a mass market, the size of the potential market that is immediately accessible is important. Compared to the expected benefits of releasing a new application, content or service on an unrestricted Internet, the expected benefits of releasing an application to those users on unrestricted Internet service offerings and, potentially, motivating users currently on restricted offerings to switch to unrestricted offerings to use the new application are much smaller.

5. BEREC should indicate clearly which forms of Quality of Service would be allowed under the guidelines. In particular, network providers should not be allowed to offer different types of service to different provider-defined classes of applications, regardless of whether the network provider treats like traffic alike. By contrast, forms of QoS should be allowed if they meet the following conditions:
(1) the different classes of service are made available equally to all applications and classes of applications;
(2) the user is able to choose whether and when to use which class of service; AND
(3) the network provider charges only its own Internet service customers for the use of the different classes of service. (This restriction would not constrain interconnection agreements in any way. Thus, payments among interconnecting networks would remain possible.)

The network neutrality debate is often framed as a debate for or against Quality of Service. As my attached paper shows, the reality is much more nuanced. 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.

Underlying the differences between the proposals are disagreements over the social benefits and costs of the different forms of Quality of Service. In this respect, the paper offers interesting new insights. Most network neutrality proponents agree that allowing network providers to offer Quality of Service exclusively to one or more applications within a class of "like" applications should be prohibited, and my paper shares that view:[31] 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

[30] See, e.g., Anderson (2006).
[31] On this form of Quality of Service, see Section "Ban Discrimination Among Like Applications and Classes of Applications" in the attached paper.

among applications within a class, which is exactly what network neutrality rules are designed to prevent.

By contrast, some participants in the debate 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."[32]

The positive stance towards forms of Quality of Service that provide like treatment is based on the assumption that discriminating among classes of applications that are not alike is socially harmless and should therefore be allowed. As this paper shows, this assumption is not correct. In many cases, discrimination among classes of applications hurts some classes of applications, even if the classes are not alike. For example, some Internet applications such as Internet telephony applications, Internet messaging applications or Internet video offerings compete with network-provider services that are sold separately from Internet access and do not run over the Internet-access portion of the network provider's access network. In these cases, discriminating against all applications in that class allows the network provider to favor its own offering without discriminating among applications within the class. Moreover, applications in a class can be harmed by differential treatment even if they do not compete directly with applications in other classes that are treated more favorably.

In addition, like treatment negatively affects several of the factors that have fostered application innovation in the past. *First*, like treatment removes the *application-blindness* 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. 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

[32] On this form of Quality of Service, see Section "Allow Discrimination Among Classes of Applications That Are Not Alike" in the attached paper.

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 where network providers make different types of service available equally to all applications and classes of applications and where users choose whether and when to use which type of service do not raise similar concerns.[33] *First*, they preserve the *application-blindness* of the network: The provision of Quality of Service is not dependent on which applications users are using, but on the Quality-of-Service-related choices that users make; thus, the network providers does not need to know anything about which applications are using its network in order for this scheme to work. The network provider only makes different classes of service available, but does not have any role in deciding which application gets which Quality of Service; this choice is for users to make. As a result, network providers cannot use the provision of Quality of Service as a mechanism to distort competition among applications or classes of applications. *Second*, since users choose when and for which applications to use which type of service (in line with the principle of *user choice*), they can get exactly the Quality of Service that meets their preferences, even if these preferences differ across users or (for a single user) over time. *Third*, in line with the principle of *innovation without permission*, an innovator does not need support from the network provider in order for his application to get the Quality of Service it needs. The only actors who need to be convinced that the application needs Quality of Service are the innovator, who needs to communicate this to the user, and the user, who wants to use the application. This greatly increases the chance that an application can get the type of service it needs.

In sum, this type of user-controlled Quality of Service offers the same potential social benefits as other, discriminatory or provider-controlled forms of Quality of Service without the social costs. With appropriate restrictions on charging and with provisions that protect the quality of the baseline service from dropping below unacceptable levels, this type of Quality of Service should be allowed under a network neutrality regime.

Opponents of network neutrality regulation have created the impression that policy makers need to choose between protecting users and application innovators against interference from network providers on the one hand and innovation in the network and the needs of network providers on the other hand. As the paper shows, it is possible to protect users and innovators while giving network providers the tools they need to manage their networks and allowing the network to evolve. Thus, regulators can have their cake and eat it, too.

[33] On this type of Quality of Service, see Section "Ban Application-Specific Discrimination, Allow Application-Agnostic Discrimination," Subsection "Allowing the Network to Evolve" in the attached paper.

6. The threat of discrimination not only affects innovation in applications, content and services in the future. The threat of discrimination affects innovation today.

BEREC's draft report on competition effects and the draft guidelines on network neutrality and Quality of Service acknowledge that discriminatory conduct may have a long-term effect on innovation in applications, content and services. However, as many conversations I have had with innovators and investors indicate, the lack of certainty about regulatory protections coupled with a very real threat of discrimination affects innovation today. The attached letter from the online video company Zediva, filed with the US Federal Communications Commission, illustrates this threat.[34] It also explains – in a very concrete way – how different forms of network neutrality rules may affect specific companies' ability to operate successfully on the Internet. As I explained elsewhere, this is just one example of many.[35] Over the past few years, many entrepreneurs have told me that potential investors identified the risk of blocking or discrimination as one of the main risks associated with their company and used this fact to justify their decision not to fund them.[36] Even those who haven't had similar conversations with funders yet are usually concerned about the problems described by Zediva. Thus, Zediva's story is not an outlier. It stands for the problems faced by many start-ups and innovators.

You may wonder why we don't hear more from entrepreneurs, if this is the case. My conversations with entrepreneurs suggest a number of reasons:

First, entrepreneurs focus on getting their product to market and making it the best product they can. They do not have the time to follow the latest twists and turns of the policy debate in Washington, Brussels or elsewhere and write letters to the relevant regulatory agencies.

Second, many do not come forward because they fear that network providers may retaliate against them in the future. I used to hear this a lot from application and service providers in the mobile space. But over the past year, this concern has started to come up in many conversations with innovators whose applications and services run over wireline networks.

Third, many start-ups do not want to draw public attention to their vulnerabilities, fearing it may scare potential investors away.

And finally, having been declined funding is not something that entrepreneurs like to brag about.

[34] Srinivasan & Gupta (2010).
[35] van Schewick (2010d).
[36] I described the experience of one start up in testimony to the Federal Communications Commission (van Schewick (2008), attached).

7. There is a gap between network providers' private interests and the public interest in open networks that BEREC's draft documents do not acknowledge.

BEREC's draft document on competition issues and its guidelines on network neutrality and Quality of Service argue that network providers will not restrict applications, content and services in socially harmful ways since more applications, content and services make their networks more attractive. This argument is only partly correct. Neither the interests of the network provider and users nor the interests of the network provider and the public are aligned. Network providers' interests often differ from users' interests, and even if they do not, network providers do not know what exactly users want.[37] Network providers' private interests and the public interests with respect to the evolution of the Internet diverge as well:[38] For a variety of reasons, network providers capture only a small part of the social value resulting from an open Internet. For example, they only capture some of the social benefits associated with application innovation or of the social benefits resulting from improved democratic discourse.[39] 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.[40] Thus, when network providers decide whether to discriminate among applications or classes of applications, the immediate private benefits of discriminating (i.e. the higher profits resulting from exclusionary conduct or from discriminatory network management) will often be higher than network providers' hyperbolically discounted share of the private benefits of refraining from discriminatory conduct. Thus, the public's and network providers' interests diverge, so regulators cannot rely on network providers to protect the public interest in an open, unrestricted Internet.

***

Thank you for considering these views. Please do not hesitate to contact me at schewick@stanford.edu if you would like to discuss these issues further.

Sincerely,
Barbara van Schewick
Associate Professor of Law and (by Courtesy) Electrical Engineering
Helen L. Crocker Faculty Scholar
Director, Center for Internet and Society
Stanford Law School

[37] See footnotes 352 to 356 in the attached paper and accompanying text.
[38] For a detailed discussion, see van Schewick (2010a), pp. 355-371 (describing the public interest), 371-375 (describing network providers' private interests and why they diverge from the public interest).
[39] van Schewick (2010a), pp. 373-374. See also Frischmann (2005); Frischmann & van Schewick (2007), pp. 400-403, 424-425; Hogendorn (2012).
[40] van Schewick (2010a), pp. 374-375.

Bibliography

Anderson, Chris. 2006. *The Long Tail: Why The Future of Business is Selling Less of More*. New York, NY: Hyperion.

Balkin, Jack M. 2009. *Testimony before the Federal Communications Commision at its Workshop on Speech, Democratic Engagement, and the Open Internet.* Federal Communications Commission. http://www.fcc.gov/events/workshop-speech-democratic-engagement-and-open-internet

Benkler, Yochai. 2000. "From Consumers to Users: Shifting the Deeper Structures of Regulation Towards Sustainable Commons and User Access." *Federal Communications Law Journal*, 52(3): 561-579.

Borthwick, John. 2010. "Neutrality or Bust." *Techcrunch.* December 19. http://techcrunch.com/2010/12/19/neutrality/

Botanical Society of America. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 10. http://apps.fcc.gov/ecfs/document/view?id=7020922879

Burnham, Brad. 2010. "Internet Access Should Be Application-Agnostic." *The Huffington Post.* December 20. http://www.huffingtonpost.com/brad-burnham/internet-access-should-be_b_799028.html

Cerf, Vinton G. 2006. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: Network Neutrality.* commerce.senate.gov/pdf/cerf-020706.pdf

Frischmann, Brett M. 2005. "An Economic Theory of Infrastructure and Commons Management." *Minnesota Law Review*, 89 (April): 917-1030.

Frischmann, Brett M. & Barbara van Schewick. 2007. "Network Neutrality and the Economics of an Information Superhighway: A Reply to Professor Yoo." *Jurimetrics Journal*, 47(4): 383–428.

Hogendorn, Christiaan. 2012. "Spillovers and Network Neutrality." In *Regulation and the Performance of Communication and Information Networks* eds. Gerald R. Faulhaber, Gary Madden & Jeffrey Petchey, 191-208. Cheltenham Glos, UK: Edward Elgar Publishing.

Lemley, Mark A. & Lawrence Lessig. 1999. Ex Parte to Federal Communications Commission. CS Dkt. No. 99-251. November 10. http://fjallfoss.fcc.gov/ecfs/document/view?id=6009850930

Lessig, Lawrence. 2001. *The Future of Ideas: The Fate of the Commons in a Connected World.* New York, NY: Vintage Books.

Lessig, Lawrence. 2002. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: The Government's Role in Promoting the Future of Telecommunications Industry and Broadband Deployment.* http://commerce.senate.gov/pdf/lessig.pdf

Lessig, Lawrence. 2006. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: Network Neutrality.* www.lessig.org/blog/archives/Lessig_Testimony_2.pdf

Lessig, Lawrence. 2008. *Testimony before the United States Senate, Committee on Commerce, Science, and Transportation, at its Hearing on: The Future of the Internet.* http://commerce.senate.gov/public/?a=Files.Serve&File_id=da3dac3c-6a90-4467-96fe-2abbeb78456b

North American Benthological Society. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 13. http://apps.fcc.gov/ecfs/document/view?id=7020923349

NYSERNet. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 13. http://apps.fcc.gov/ecfs/document/view?id=7020923371

Reed, David P. 2010. "A Respone to Barbara van Schewick: Code Needs (Only a Little) Help From the Law." *dpr.* December 15. http://www.reed.com/blog-dpr/?p=85

Srinivasan, Venky & Vivek Gupta. 2010. Ex Parte Letter to Federal Communications Commission. GN Docket No. 09-191. December 10. http://apps.fcc.gov/ecfs/document/view?id=7020923207

The Council of Scientific Society Presidents. 2010. Ex Parte Letter to Federal Communications Commission. WC Dkt. No. 07-52. December 13. http://apps.fcc.gov/ecfs/document/view?id=7020923181

van Schewick, Barbara. 2004. "Architecture and Innovation: The Role of the End-to-End Arguments in the Original Internet." PhD Dissertation. Technical University Berlin.

van Schewick, Barbara. 2007. "Towards an Economic Framework for Network Neutrality Regulation." *Journal on Telecommunications and High Technology Law*, 5(2): 329-391.

van Schewick, Barbara. 2008. *Oral Testimony at the Federal Communications Commission Second En Banc Hearing on Broadband Management Practices.* Federal Communications Commission. www.fcc.gov/broadband_network_management/041708/vanschewick-oral.pdf

van Schewick, Barbara. 2010a. *Internet Architecture and Innovation.* Cambridge, MA: MIT Press.

van Schewick, Barbara. 2010b. *Opening Statement at the Federal Communications Commission's Workshop on Approaches to Preserving the Open Internet.* Federal Communications Commission. http://www.law.stanford.edu/display/images/dynamic/publications_pdf/schewick-statement-20100428.pdf

van Schewick, Barbara. 2010c. *Opening Statement at the Federal Communications Commission's Workshop on Innovation, Investment and the Open Internet.* Federal Communications Commission. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020382126

van Schewick, Barbara. 2010d. "Start-Up Video Company Asks FCC to Improve Open Internet Proposal." *Internet Architecture and Innovation Blog.* December 12. http://netarchitecture.org/2010/12/start-up-video-company-files-concerns-about-fcc-open-internet-proposal/

van Schewick, Barbara. 2012. *Network Neutrality: What a Non-Discrimination Rule Should Look Like.* Center for Internet and Society White Paper. http://cyberlaw.stanford.edu/publications/network-neutrality-and-quality-service-what-non-discrimination-rule-should-look

Wilson, Fred. 2010. "Regulation Strangulation." *AVC.* August 12. http://www.avc.com/a_vc/2010/08/regulation-strangulation.html

Wu, Tim. 2003. "Network Neutrality and Broadband Discrimination." *Journal on Telecommunications & High Technology Law*, 2: 141-175.

Wu, Tim. 2004. "The Broadband Debate: A User's Guide." *Journal on Telecommunications & High Technology Law*, 3(1): 69-95.

Wu, Tim & Lawrence Lessig. 2003. Ex Parte Submission to Federal Communications Commission. CS DKT. No. 02-52. August 22. http://fjallfoss.fcc.gov/ecfs/document/view?id=6514683885