Network Neutrality and Quality of Service

What a Non-Discrimination Rule Should Look Like

Barbara van Schewick | June 11, 2012

EXECUTIVE SUMMARY

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, including the United States, the European Union, the United Kingdom, France and Germany, have been investigating whether legislative or regulatory action is needed to limit the ability of providers of Internet access services[1] to interfere with the applications, content and services on their networks.

But network neutrality proposals are not identical, and cover a range of potential rules. Policy makers considering whether to enact network neutrality rules need to answer a series of questions in order to decide which, if any, network neutrality rules they should adopt. The first question is: "Do we need a rule against blocking, i.e. a rule that forbids network providers from blocking applications, content and services on their networks?" A rule against blocking is part of all network neutrality proposals; it is the one rule on which all network neutrality proponents agree. This paper assumes that the case for a rule against blocking has been made.

This paper focuses on the second question: "Should the rules also ban discrimination (i.e., differential treatment that falls short of blocking), and, if yes, which forms of differential treatment should be banned?" For example, if a network provider slows down Internet video applications like Hulu or YouTube that compete with the network provider's own Internet video application, or provides quality of service only to its own Internet video application, should these practices be prohibited, too?

The answer depends in part on what the goals of network neutrality regulation should be. According to most network neutrality proponents, 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. In particular, network neutrality rules aim to foster innovation in applications, protect users' ability to choose how they want to use the network, without interference from network providers, and preserve the Internet's ability to improve democratic discourse, facilitate political organization and action and to provide a decentralized environment for social, cultural and political interaction in which anyone can participate. Thus, network neutrality proponents base their calls for regulation on a theoretical framework that considers a broad range of economic and non-economic harms. The Federal Communications Commission's Open Internet rules are based on this broader framework as well. By contrast, some participants in the debate view the network neutrality debate through an antitrust lens and evaluate and address concerns about blocking, discrimination or other practices based on an antitrust framework. Among network neutrality proponents, this is a minority position. I address the debate over the right framework elsewhere.[2] This paper assumes that the case for the broader theoretical framework has been made.

---

[1] Throughout this paper, I use the terms "providers of Internet access service", "Internet service providers", or "network providers" interchangeably.
[2] van Schewick (2012).

The decision for a non-discrimination rule has important implications: Non-discrimination rules affect how the core of the network can evolve, how network providers can manage their networks, and whether they can offer Quality of Service (QoS). While the original Internet provides a single best-effort service for all packets (i.e., the network does its best to deliver data packets, but does not provide any guarantees with respect to delay, bandwidth or losses), a network that provides Quality of Service offers different types of service to different data packets. 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. Different applications have different requirements with respect to reliability, bandwidth or delay. 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. Thus, whether network providers are able to offer Quality of Service may have implications for the types of applications that the Internet can support.

Answering these questions is the goal of this paper. It sets out a framework for evaluating network neutrality rules and uses this framework to evaluate existing proposals for a non-discrimination rule and the non-discrimination rule adopted by the Federal Communications Commission in its Open Internet Order. In the process, it explains how the different non-discrimination rules affect network providers' ability to offer Quality of Service and which, if any, forms of Quality of Service a non-discrimination rule should allow.

A Framework for Evaluating Network Neutrality Rules

Network neutrality proponents generally agree that network neutrality rules should preserve the Internet's ability to serve as an open, general-purpose infrastructure that provides value to society over time in various economic and non-economic ways. 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 choose among different options for network neutrality rules. As the paper shows, this framework can also be used 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.

Most generally, any network neutrality rule should realize the goals of network neutrality regulation, while minimizing the social costs. Thus, policy makers should look for a rule that fosters application innovation, protects user choice and preserves the Internet's social, cultural and political potential while avoiding unnecessary social costs. These general requirements can be translated into the following framework:

Box 1
Framework for Evaluating Network Neutrality Rules

- A network neutrality rule should meet the following criteria:

- It should preserve the factors that have allowed the Internet to serve as a platform for application innovation, free speech and decentralized economic, social, cultural and political interaction in the past:

  - *User choice:* Users independently choose which applications they want to use, without interference from network providers.
  - *Innovation without permission:* Innovators independently choose which applications they want to pursue; they do not need support or "permission" from network providers in order to realize their ideas for an application.
  - *Application-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.
  - *Low costs of application innovation:* The costs of application innovation are low.

- It should not constrain the evolution of the network more than is necessary to reach the goals of network neutrality regulation.

- It should make it easy to determine which behavior is and is not allowed to provide much-needed certainty for industry participants.

- It should keep the costs of regulation low.

In its Open Internet order, the FCC adopted a similar framework. Preserving user choice and innovation without permission is an explicit purpose of the rules. Thus, these factors can be used to guide the interpretation of individual provisions of the Open Internet rules. Moreover, the text of the order explicitly specifies that the FCC will interpret key provisions of its rules – the non-discrimination rule for fixed broadband Internet access and the reasonable-network-management exception – based on how well they preserve three of the factors used to evaluate alternative options for non-discrimination rules and specific discriminatory conduct throughout this paper: user choice, application-blindness and innovation without permission.[3] Thus, the results of this paper indicate how the non-discrimination 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.

---

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

Evaluating Proposals for Non-Discrimination Rules

Scope of Non-Discrimination Rules

Non-discrimination rules apply to any form of differential treatment that falls short of blocking. The most obvious examples involve differential handling of data packets associated with different applications or uses. For example, non-discrimination rules decide whether a network provider is allowed to provide a low-delay service only to its own streaming video application, but not to competing streaming video applications.

The scope of non-discrimination rules is, however, not restricted to differential handling of packets in the network. Network neutrality rules aim at preventing network providers from distorting the playing field among applications[4] or classes of applications and from interfering with users' choices regarding the use of the network. In line with this goal, non-discrimination rules apply to any form of differential treatment that may make some applications or classes of applications or uses relatively more attractive to users than others. Thus, they cover, for example, network provider practices that count only some applications, but not others towards a user's monthly bandwidth cap, or pricing plans that impose different bandwidth-adjusted Internet access charges on different applications.

All-or-Nothing Approaches

A first set of proposals takes an all-or-nothing approach to discrimination.

The first approach would *allow all forms of discrimination*, which is equivalent to not adopting a non-discrimination 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.

As the paper shows, both approaches are flawed. Banning all discrimination is over-inclusive and restricts the evolution of the network more than necessary to protect the values that network neutrality rules are designed to protect. Allowing all discrimination is under-inclusive and effectively makes the rule against blocking meaningless.

---

[4] Throughout this paper, I use "applications" as shorthand for "applications, content, services and uses."

Case-by-Case Approaches

A second set of proposals 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 discrimination in the future. Whether certain discriminatory conduct meets these criteria would be determined by the regulatory agency in future case-by-case adjudications.

The paper discusses three approaches in this category:

- A rule that *bans discrimination that would be considered harmful under an antitrust standard*.
- A rule that *bans discrimination that is anticompetitive or harms users*.
- A rule that *bans discrimination that is unreasonable*.

Whether and under which conditions these rules would allow some and, if yes, which forms of Quality of Service would have to be decided in future case-by-case adjudications.

These approaches raise two sets of concerns: First, the first and second approach fail to capture many instances of discrimination that network neutrality proponents are concerned about. Network neutrality proponents often think of discriminatory conduct that favors an application over others as a distortion of competition and, therefore, as "anticompetitive," and assume such behavior would be captured by an antitrust framework. This assumption is not correct. As the paper shows, the term "anticompetitive" has a much narrower scope in antitrust law than an intuitive interpretation of the term would suggest. More generally, while an antitrust framework focuses on a narrow set of economic harms, the theoretical framework underlying calls for network neutrality regulation addresses a broader range of economic and non-economic harms. As a result, rules that ban behavior that is anticompetitive or violates an antitrust framework would often classify differential treatment as socially beneficial that network neutrality proponents would consider socially harmful, making it impossible to successfully challenge behavior that network neutrality are concerned about.

Second, these proposals 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. 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, case-by-case approaches create high costs of regulation and tilt 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 discriminatory conduct in individual

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

In spite of the considerable social costs, the strategic incentives of regulators or legislators and of the big stakeholders on both sides of the network neutrality debate are aligned in favor of such a scheme. 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 two big stakeholders (Verizon and Google) on opposite sides of the debate.

More Nuanced Rules

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

Formal Approaches

The first approach in this group *bans discrimination that is not disclosed*, distinguishing between socially beneficial and socially harmful practices using a formal criterion. This rule would allow all forms of Quality of Service as long as they are disclosed. This approach is based on the idea that if a network provider discriminates against an application that users would like to use, users can switch to another network provider who does not discriminate against the affected application. The threat of switching, proponents of this approach assume, will discipline providers.

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

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 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, even if there is competition in the market for Internet access services, disclosure cannot replace substantive regulation as a tool to discipline providers.

Substantive Approaches

The second set of approaches in this group relies on substantive criteria to specify in advance which forms of differential treatment are allowed.

The first approach – which is the approach favored by this paper – bans application-specific discrimination, but allows application-agnostic discrimination. (Again, this paper uses "applications as shorthand for "applications, content, services, and uses.") 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.[5] The rule would 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.

Under this rule, a network provider would not be allowed to treat Vonage differently from Skype, YouTube differently from Hulu, or the website of the New York Times differently from the website of the Wall Street Journal or Free Press. That would be discrimination based on application. Nor would it be allowed to treat online video differently from e-mail, treat applications that use the BitTorrent protocol differently from applications that do not use this protocol, or treat latency-sensitive applications differently from latency-insensitive applications. That would be discrimination based on class of application. But it would be allowed to treat data packets differently based on criteria that have nothing to do with the application or class of application. For example, it could give one person a larger share of the available bandwidth if that person has paid for a higher tier of Internet service (e.g., if that person has paid for the "Up to 6 Mbps" Internet service packet instead of the "Up to 3 Mbps" Internet service packet).

This approach strikes the best balance between social benefits and social costs: 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

---

[5] Relevant characteristics of an application include what this 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 the rule.

Internet's ability to reach its social, cultural or political potential. Based on these insights, the rule 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 rule 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.[6] 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.[7]

Application-agnostic network management coupled with user-controlled prioritization[8] 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 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. [9,10,11]

---

[6] That would be application-agnostic discrimination.

[7] That would be discrimination based on application or class of application.

[8] 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.

[9] See also the discussion on of user-controlled Quality of Service on pp. xiii-xiii of this Executive Summary below.

[10] This restriction would not constrain interconnection agreements in any way. Thus, payments among interconnecting networks would remain possible.

[11] 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.

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.

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

The second approach in this group *bans discrimination among like applications and classes of applications, but allows discrimination among classes of applications that are not alike and application-agnostic discrimination*. Like the approach favored by this paper (and for the same reasons), this rule bans discrimination among like applications or among like classes of applications and allows application-agnostic discrimination.[12] Contrary to that approach, however, this rule would allow network providers to treat classes of applications differently as long as they treat like traffic alike. This requirement is also called "like treatment."

Under this rule, a network provider would be allowed to offer different types of service to different provider-defined classes of applications that are not alike, as long as it does not discriminate among classes of applications that are alike or discriminate among like applications within a class. For example, a network provider would be allowed to treat Internet telephony differently from e-mail (e.g., by providing a low-delay service to Internet telephony, but not to e-mail), but it would not be allowed to treat Vonage differently from Skype, or Gmail differently from Hotmail.

This rule is based on the assumption that discrimination among classes of applications that are not alike is socially harmless, and should therefore be allowed. As the paper shows, this assumption is not correct.[13] Like treatment creates considerable social costs. In particular, like treatment 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. As a result, the rule's social costs may be similar to the social costs of the case-by-case approaches described above.

---

[12] The approach proposed by this paper bans any discrimination among applications or classes of applications that is based on application-specific criteria, regardless of whether the applications or classes of applications are alike or not.

[13] For a summary of the argument, see the discussion of like treatment on pp. xii-xiii of this Executive Summary below.

The Federal Communications Commission's Non-Discrimination Rule and its Impact on Quality of Service

In December 2010, the Federal Communications Commission adopted network neutrality rules. The rules went into effect in November 2011. They include a non-discrimination rule that *bans discrimination that is "unreasonable."* Whether a certain discriminatory conduct meets these criteria will be determined by the FCC in case-by case-adjudications. The rule applies only to fixed, not to mobile broadband Internet access service and has an exception for reasonable network management.

The text of the order specifies the factors that the FCC will use to determine whether specific discriminatory conduct is "unreasonable:" transparency (i.e. whether the discrimination is disclosed), end user-control, use-agnosticm, and the conformity of a practice with standards and best practices. The same factors will guide the evaluation of network management practices in the context of the reasonable-network-management exception.

Use-agnostic discrimination (or "application-agnostic" discrimination), the order explains, is differential treatment that does not discriminate among specific uses of the network or among classes of uses. According to the order, use-agnostic discrimination is likely to be reasonable, which suggests, in turn, that differential treatment that discriminates among specific uses of the network or classes of uses is likely to be unreasonable. This is the same substantive standard as the one used by the non-discrimination rule proposed by this paper.

Allowing use-agnostic discrimination, but banning discrimination among specific uses or classes of uses preserves the application-blindness of the network. Thus, in evaluating whether discriminatory conduct is reasonable, the FCC will consider how well the conduct preserves two of the four factors – application-blindness and user choice – that have fostered application innovation in the past.[14] In addition, the first section of the Open Internet rules lists preserving end user control, user choice and the freedom to innovate without permission – the third factor that has fostered application-innovation in the past – as explicit purposes of the rules. As explicit purposes of the Open Internet rules, these factors can be used to interpret any provision of these rules, including the non-discrimination rule.

As this paper shows, using these factors – application-blindness, user choice and innovation without permission – as guidelines for evaluating behavior provides clear answers regarding which types of differential treatment should and should not be allowed. For example, while the order does not discuss how the different forms of Quality of Service discussed in this paper would be evaluated under the "no unreasonable discrimination" standard, the analysis in this paper suggests which results the FCC will reach, if it takes these factors seriously.[15] In particular, under the FCC's standard as clarified by the order, user-controlled Quality of Service that meets the conditions outlined above is likely to be reasonable.[16] By contrast, offering

---

[14] The other factors are innovation without permission and low costs of application innovation.

[15] For a more detailed summary of the argument, see the Section "Network Neutrality and Quality of Service" in this Executive Summary below.

[16] The conditions are described on p. viii of this Executive Summary above. According the Open Internet Order, charging application or content providers who are not the network provider's Internet access service customers for

Quality of Service exclusively to one or more applications within a class of "like" applications is unlikely to be reasonable. 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 like applications within a class).

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. In particular, the order explicitly rejects the view that the rules should only prohibit discrimination that is "anticompetitive." This is an important clarification. Such a rule (or an interpretation of the FCC's rule that restricted unreasonable discrimination to discrimination that is anticompetitive) would have made it impossible to successfully bring complaints against many types of discriminatory conduct that network neutrality proponents are concerned about.

Overall, the FCC's non-discrimination rule (as clarified by the order) provides more clarity to industry participants than the case-by-case approaches discussed above and provides additional guidance to the bureaus within the FCC which may enforce the order. Substantively, the framework adopted by the order reinforces key values that were at the core of the Internet's success. Still, compared with the approach proposed by this paper, the rule provides less certainty to market participants and leaves considerably more discretion to later adjudicators. As a result, the rule still shares many of the social costs with the less precise case-by-case approaches discussed above.

Network Neutrality and Quality of Service

The network neutrality debate is often framed as a debate for or against Quality of Service. This is not correct. Some proposals take an all-or-nothing approach to discrimination. They ban or allow all forms of discrimination and, consequently, Quality of Service. Most proposals take a more nuanced position. They allow some, but not all forms of Quality of Service, with different proposals drawing the line between acceptable and unacceptable forms of Quality of Service in different ways.

Often, it is not immediately apparent how a specific non-discrimination rule affects network providers' ability to offer Quality of Service. To address this problem, this paper explores the effect of the various proposals on the different forms of Quality of Service. The results of this analysis are summarized in Table 1: The Impact of Different Non-Discrimination Rules on Quality of Service below.

Underlying the differences between the proposals are disagreements over the social benefits and costs of the different forms of Quality of Service. In this respect, the paper offers

---

prioritized or otherwise enhanced access to its Internet access service customers is unlikely to be reasonable under the order's non-discrimination rule. Thus, under the Open Internet rules, network providers are (most likely) allowed to charge only their own Internet access service customers for any differential treatment allowed by the non-discrimination rule. See van Schewick (2010b).

interesting new insights. Most network neutrality proponents agree that allowing network providers to offer Quality of Service exclusively to one or more applications within a class of "like" applications should be prohibited, and this paper shares that view:[17] 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."[18]

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

---

[17] On this form of Quality of Service, see Section "Ban Discrimination Among Like Applications and Classes of Applications" in the paper.
[18] On this form of Quality of Service, see Section "Allow Discrimination Among Classes of Applications That Are Not Alike" in the paper.

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 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.[19] *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.

---

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

Network Neutrality and Charging for Quality of Service

If policy makers adopt a non-discrimination rule that allows network providers to offer some form of Quality of Service, they need to decide whether, and if yes, whom network providers should be allowed to charge for it. Again, policy makers have a number of options, each supported by at least some proponents of network neutrality: (1) the network provider is not allowed to charge anyone for the use of Quality of Service (though it can increase the general price for Internet service); (2) it can charge only its Internet service customers; (3) it can charge its Internet service customers and/or application and content providers, but is required to offer the service to application and content providers on a non-discriminatory basis; (4) it can charge its Internet service customers and/or application and content providers. Concerns about offering differential treatment and about charging for it are driven by different sets of policy considerations which should be considered and evaluated separately. I take up restrictions on charging elsewhere.[20] There, I argue that network providers should only be allowed to charge their own Internet service customers for any differential treatment allowed by the non-discrimination rule.[21]

Finally, a network provider who is allowed to charge for Quality of Service has an incentive to degrade the quality of the baseline, best-effort service to motivate users to pay for an enhanced type of service. To mitigate this problem, any network neutrality regime that allows network providers to charge for Quality of Service should require the regulatory agency in charge of enforcing the network neutrality rules to monitor the quality of the baseline service and set minimum quality standards, if the quality of the baseline service drops below appropriate levels.[22]

---

[20] van Schewick (2010d); van Schewick (2010e); van Schewick (2010c), pp. 278-280, 290-293

[21] This restriction would not constrain interconnection agreements in any way. Thus, payments among interconnecting networks would remain possible.

[22] For a discussion of this requirement, see van Schewick (2010a), pp. 10-11. The European Union has adopted a similar rule following its review of the regulatory framework for telecommunications services. See Article 22(3) of the Universal Service Directive; European Commission (2007), pp. 92, 95-97, 101.

Table 1

The Impact of Different Non-Discrimination Rules on Quality of Service

Forms of QoS / Rules	Provider-controlled QoS to individual applications within a class of like applications	QoS to provider-defined classes of applications	User-controlled QoS
Allow all discrimination	allowed*	allowed*	allowed*
Case-by-case approaches	unclear	unclear	unclear
Ban discrimination that is not disclosed	allowed, if disclosed*	allowed, if disclosed*	allowed, if disclosed*
Ban discrimination among like applications and classes of applications, allow discrimination among classes of applications that are not alike and application-agnostic discrimination ("like treatment")	banned	allowed, as long as like traffic is treated alike*	allowed, if (1) different classes of service are made available equally to all applications and classes of applications; AND (2) user is able to choose whether and when to use which class of service.*
Ban application-specific discrimination, allow application-agnostic discrimination	banned	banned	allowed, if (1) different classes of service are made available equally to all applications and classes of applications; (2) user is able to choose whether and when to use which class of service.*
FCC's non-discrimination rule	likely banned	likely banned	likely allowed, if (1) different classes of service are made available equally to all applications and classes of applications; (2) user is able to choose whether and when to use which class of service; AND (3) network provider charges only its own Internet service customers for the use of the different classes of service.**
Ban all discrimination	banned	banned	banned

* If policy makers adopt a non-discrimination rule that allows network providers to offer some form of Quality of Service, they need to decide whether, and if yes, whom network providers should be allowed to charge for it. As I argue elsewhere, network providers should only be allowed to charge their own Internet service customers for any differential treatment allowed by the non-discrimination rule.

** According the FCC's Open Internet Order, charging application or content providers who are not the network provider's Internet access service customers for prioritized or otherwise enhanced access to its Internet access service customers is likely to be unreasonable. Thus, under the Open Internet rules, network providers are allowed to charge only their own Internet access service customers for any differential treatment allowed by the non-discrimination rule.

REFERENCES

European Commission. 2007. "Commission Staff Working Document. Impact Assessment". SEC(2007) 1472. http://ec.europa.eu/information_society/policy/ecomm/doc/library/proposals/1472/comm_pdf_sec_2007_1472_1_en_documentdetravail.pdf.pdf

van Schewick, Barbara. 2010a. *Background Paper for the Federal Communications Commission's Workshop on Approaches to Preserving the Open Internet in Seattle, WA on April 28, 2010.*

van Schewick, Barbara. 2010b. "The FCC's Open Internet Rules – Stronger than You Think." *Internet Architecture and Innovation Blog.* December 27. http://netarchitecture.org/2010/12/the-fcc's-open-internet-rules-–-stronger-than-you-think/

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

van Schewick, Barbara. 2010d. *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. 2010e. *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. 2012. *Should Network Neutrality Learn from Antitrust?* Unpublished Working Paper

About the Author

Barbara van Schewick is an Associate Professor of Law and Helen L. Crocker Faculty Scholar at Stanford Law School, an Associate Professor (by courtesy) of Electrical Engineering in Stanford University's Department of Electrical Engineering, Director of Stanford Law School's Center for Internet and Society, and a leading expert on network neutrality.

Her research on the economic, regulatory, and strategic implications of communication networks bridges law, networking and economics. Her book Internet Architecture and Innovation (MIT Press 2010) is considered to be the seminal work on the science, economics and policy of network neutrality. Her papers on network neutrality have influenced regulatory debates in the United States, Canada and Europe. Van Schewick has testified before the FCC in en banc hearings and official workshops. In October 2010, van Schewick received the Research Prize Technical Communication 2010 from the Alcatel-Lucent Stiftung for Communications Research for her pioneering work in the area of Internet architecture, innovation and regulation.

About the Center for Internet and Society

The Center for Internet and Society (CIS) is a public interest technology law and policy program and a part of the Law, Science and Technology Program. The Center brings together academics from a variety of disciplines with policy makers and a wide range of stakeholders to study the interaction of new technologies and the law. The Center strives to improve both technology and law, encouraging decision makers to design both as a means to further democratic values. Along with conducting research and policy analysis, the Center sponsors legal fellowships, organizes events to foster discussion of critical policy issues, and provides educational opportunities for law students to conduct applicable research and policy analysis in this field.