NETWORK NEUTRALITY IN AN INCREASINGLY DIVERSE WORLD

By Christopher S. Yoo[1]

I. INTRODUCTION

Over the course of the last year, regulatory interest in network neutrality has intensified. In the U.S., the Federal Communications Commission ("FCC") responded to the judicial invalidation of its initial effort to regulate network neutrality by adopting a second Open Internet Order ("OIO") in February 2015, which was upheld by the courts in June 2016. In Europe, the European Parliament's inclusion of strong network neutrality provisions in the Single Telecom Market legislation prompted opposition from the Council of the European Union, which ultimately led to the adoption of a compromise solution in October 2015. The Body of European Regulators for Electronic Communications ("BEREC") issued guidelines in August 2016 regarding the best way to implement this legislation. In India, the Telecommunications Regulatory Authority of India ("TRAI") adopted a regulation in February 2016 that prohibits all discriminatory tariffs for data services on the basis of content.

The basic premise behind network neutrality is that providers should not favor any network traffic based on its source or the content or application with which it is associated, either by blocking certain traffic or by charging differential amounts for delivering it. Network neutrality is animated by a vision in which the Internet adheres to the best-efforts architecture around which the network was initially organized and that has to date proven so successful in promoting innovation. In essence, this vision effectively requires network

[1] Christopher S. Yoo is the John H. Chestnut Professor of Law, Communication, and Computer & Information Science and the Founding Director of the Center for Technology, Innovation and Competition at the University of Pennsylvania.

providers to focus their efforts on providing a single class of service. Network neutrality proponents claim that doing so promotes innovation by ensuring that all creators and content and applications have access to the broadest possible market and by preventing them from having to pay differential amounts for network access.

The benefits provided by having a single platform for which all content and applications providers can design their offerings unquestionably represents an important source of value on the Internet. What is often overlooked is that network size is not the only source of value. The benefits of mandating net uniformity must be traded off against other aspects, such as the inherent loss of variety and the inability to appeal to new users.

In other words, the debate has failed to provide any basis for distinguishing the circumstances under which having a single, monolithic design is optimal from those in which a more diverse architecture might serve consumers better. A key consideration is the homogeneity or heterogeneity of consumer preferences. If people's demand for network services is relatively uniform, the optimal course is to deploy a single network designed to satisfy those preferences. As demand for network services becomes more heterogeneous, eventually optimality will demand that network providers offer a broader range of services to meet that demand. This variety can both better satisfy users who already in the market and attract in new users who are not yet in the market.

This analysis suggests the need to understand the tradeoffs inherent in any decision to standardize around any particular design and to appreciate that any such standardization can have hidden costs. Only by framing standardization as a question of optimality can regulators discern when mandating network neutrality might be good for consumers and when economic welfare may be better served by permitting a greater diversity of network offerings.

II. BETTER SATISFACTION OF CURRENT USERS

Competition policy has long understood the potential benefits of standardization through the lens of the now substantial literature on network economic effects. As is now well understood, network economic effects exist when the value of network increases with the number of users connected to it.

Were network size the only relevant source of value, standardization would always be the preferred course, and all networks would offer only a single class of service. In reality, networks have always offered a variety of services. Unless those alternative services are to be dismissed as mistakes, making sense of this behavior requires a better understanding of the value created by alternative classes of service. What emerges is an appreciation of how variety can be a source of consumer welfare outside the price-quantity space. The advantages from being part of a larger network must be traded off against the disadvantages of surrendering product variety.

A. The Tradeoff from Between Standardization and Variety

Discussions of network economics often overlook the fact that size is only one of many potential sources of economic welfare in network industries. When tastes vary, consumers can also generate surplus by consuming a product more closely suited to their tastes.

The tradeoff is framed nicely by a short and easily accessible article by Joseph Farrell and Garth Saloner.[2] Suppose that there are two groups of users, each with a preference for a different standard. Each group faces a choice. It can adopt the standard preferred by the other group, which would provide it with the benefits of being part of a larger network, but would require it to forego the benefits of employing its preferred standard. Or it can adhere to its preferred standard, which would allow it to realize the benefits from using the standard more closely aligned with its preferences, but reduce the size and thus the value of the network in which it participates.

The implications of this model are elegant and clear. If preferences are uniform, variety provides no value, and consumers are best served if everyone is part of a single network. The optimal tradeoff changes as preferences become increasingly heterogeneous. At some point, the value from variety dominates the value from network size. When that occurs, the optimal outcome is for users to be part of separate networks.

B. Implications for Network Neutrality

This simple model offers some real insights into the debate over requiring networks to offer only a single, uniform class of service. The current Internet is designed to provide a particular cluster of services. The best-efforts architecture does not provide any guarantee as to how quickly any particular traffic will arrive at its destination or indeed if it will arrive at all. The applications that dominated the early Internet (email and web browsing) were fairly tolerant of delays of a fraction of a second and of irregularities in the timing with which individual packets arrive (known as "jitter"). They were able to engineer around the lack of reliability at the cost of introducing even larger delays.

The modern Internet is quite different. In the developed world, the user base has exploded both in terms of numbers and heterogeneity, and users are employing a more diverse range of applications that are making more intensive and varied. Perhaps most salient is the rise of streaming video, which according to Cisco represented 68 percent of consumer Internet traffic as of 2015 and is projected to grow to 82 percent by 2020. Sandvine reports that in 2016 two applications, Netflix and YouTube, together represented more than half of all downstream peak-period traffic, with 35 percent and 18 percent respectively.

Streaming video is just one example of how modern applications are demanding a different combination of network services. It demands significantly more bandwidth and is far more sensitive to jitter, while at the same time being more tolerant of lack of reliability as well as initial delays in initiating the video stream. Other applications, such as streaming audio and Voice over Internet Protocol ("VoIP"), require different clusters of services.

[2] Joseph Farrell and Garth Saloner, "Standardization and Variety," *Economics Letters* 20: 71, 1986.

Network Services Demanded by Established Internet Applications

Application	Bandwidth	Reliability	Delay	Jitter
E-mail	Low	High	Low	Low
Web browsing	Medium	High	Low	Low
Streaming audio	Medium	Low	Low	High
Streaming video	High	Low	Low	High
VoIP	Low	Low	High	High

Newer applications are placing still different demands on the network. Interactive video is far more sensitive than streaming video to delay. The advent of the Internet of Things is creating a new array of applications that require still different clusters of network services. And waiting in the wings are even more experimental technologies, such as virtual reality, autonomous vehicles and technologies that have yet to be conceived, that may place still different demands on the Internet.

Network Services Demanded by Emerging Internet Applications

Application	Bandwidth	Reliability	Delay	Jitter
Interactive video	High	Low	High	High
Smart metering	Low	Medium	High	Low
SCADA	High	High	Low	Low
Video surveillance	High	High	Medium	High
Mobile workforce	Low	High	Low	Low
Smart homes	Low	Medium	Low	Low

The increasing variety of demands that new applications are placing on the Internet suggests that forcing the Internet into a single class of service may not always represent the best way to promote innovation. While certainly beneficial to content and applications whose demands are aligned with the cluster of services provided by current design of the Internet, mandating net uniformity would harm innovations that require something different from the network.

As a technical matter, no single architecture can perform all of these combinations of

functions equally well. Indeed, the engineering literature is replete with articles recognizing that there are many tasks that the current Internet is not well designed to implement, including security, mobility and mass media distribution. While these aspects were less important when the Internet first emerged, modern Internet usage has made them mission critical for many users.

In short, the emergence of these new applications has increased the heterogeneity in the demands that people are placing on the Internet. As these demands become sufficiently heterogeneous, at some point, requiring that the entire network offer only a single class of service becomes suboptimal and bad for consumers. Enforcement of overly restrictive network neutrality policies would reduce the consumer surplus being generated by the network and would force those who value new uses the most to turn to private networking and non-compliant technologies, such as Multi-Protocol Label Switching ("MPLS"), to get the networking services that they need.

Network diversity can also open new dimensions along which network providers can appeal to customers. Consider T-Mobile's Binge On, which allows users to stream video without having that traffic count against their data caps. This plan is designed to appeal to young people and others who place a particularly high value on the ability to watch streaming video. The ability to design services to appeal to subsets of the overall customer base represents a new source of consumer value. In the process, it opens new dimensions along which network providers can compete.

In the developing world, the rationale for providing different levels of service plays out slightly differently. In countries where the ability to pay is more limited, offering service plans optimized for the applications that users value the most can allow providers to serve existing customers at lower cost. In this case, network diversity is designed not to address the increasing heterogeneity of applications, but rather to reflect differences in the nature and intensity of consumer demand.

The recognition that different types of users need different clusters of services is the reason that telecoms regulation has long permitted network providers to offer different classes of service so long as they make each class available to anyone who wanted it. In other words, non-discrimination required treating all users within each class of service equally; it did not forbid creating premium services and charging those who place a particularly high value on those services more.

The 2015 OIO represents a fairly significant break from these well-established principles. It explicitly bans paid prioritization, defined to include favoring some traffic over other traffic in return for compensation. While all other aspects of the Order are subject to an exception for reasonable network management, this exception does not apply to the ban on paid prioritization.

These provisions are designed to prevent network providers from creating different classes of service and charging customers different amounts for them. What has gone largely unrecognized is the extent to which they represent a significant deviation from the established principles of common carriage and telecom regulation. Although network neutrality often purports to be a return to approaches taken in the past, closer inspection underscores just how radical forcing the Internet into a single class of service actually is.

III. ATTRACTING IN NEW USERS

Allowing network providers to offer different classes of service can do more than just increase the welfare of those who are already using the Internet. It can expand surplus still further by attracting in new users. Regular Internet users may find it surprising to discover that the primary reason for non-adoption is the lack of perceived need and relevance, which ranks well above a lack of digital literacy, availability and cost as the primary barrier to adoption. Indeed, this finding is consistent in both the developing and the developed world.

Alternative service plans can help address this problem. Consider the emerging practice of zero rating, which allows users to access certain applications without having that traffic count against their data caps. The most prominent example is Facebook's "Free Basics" program, which provides users with free feature phones and free access to a suite of applications that is open to any service that can satisfy the technical requirements. Because the Internet is generally regarded as an "experience good" that must be actively used before one can appreciate its benefits, such programs can play a key role in demonstrating to non-adopters the benefits of joining the Internet community.

Some critics regard zero rating plans as network neutrality violations because they give discounted (in fact, free) access to certain applications, but not others. While generality offers real benefits to consumers, the engineering community has long recognized that it comes at a cost. In a world where many subscribers can pay no more than US$ 3 per month for Internet access, many network providers are offering service-specific plans that support only those services that customers want the most, such as email, web browsing, and a limited number of other applications. The ability to provide services that are less than the full range normally provided by the Internet can create real value by demonstrating the value of adoption to those who have not yet done so and by lowering the burden that providers must meet in order to offer service.

IV. A BRIEF COMMENT ABOUT INTERCONNECTION

There is one aspect of the FCC's 2015 OIO that represents a significant expansion of the network neutrality debate that has received almost no attention. Throughout almost the entirety of the debate, network neutrality focused exclusively on how traffic was treated within the network of last-mile Internet service providers. At some point between the 2014 proposal of the rules that would become the current OIO and the 2015 adoption of those rules, the scope of regulation was expanded to include how networks hand off traffic to each other.

Although many aspects of the debate over network neutrality seem well established, this aspect is quite new. A full analysis of the implications of this expansion exceeds the scope of this brief article. Suffice to say, regulators should exercise extreme caution before embracing such a novel development until its implications are better understood.

V. CONCLUSION

The value of a single Internet operating on a uniform set of principles is both intuitive and seductive. Closer inspection reveals that requiring Internet service providers to provide only a single class of service can have a cost both to consumers and to innovation. Mandating

network neutrality runs the risk of depriving consumers of the benefits of variety. It also foregoes the opportunity of providing those who have not yet adopted the Internet because of lack of perceived need with a low-cost way to discover its benefits. Such costs loom particularly large in a world where what consumers want from the network has become increasingly diverse and in which more than half of the world's citizens are not yet on the network.

Indeed, there is an "ever was, ever shall be" quality to the debate over network neutrality that is quite at odds with the way that both engineers and economists approach problems. Both disciplines tend to frame issues in terms of optimality that is contingent on the particular circumstances. Neither would pretend that a single solution exists that is inherently superior regardless of the underlying conditions. Instead, they tend to adopt a more dynamic, context-sensitive approach that examines whether technical or economic changes have altered the optimal outcome. The debate over network neutrality would be far healthier if it examined how patterns of usage have changed over time and created a framework for understanding the implications of those changes for the optimal network design instead of treating the existing architecture as an inviolable feature that necessarily must be preserved.