

**BEFORE THE
Federal Communications Commission
WASHINGTON, D.C.**

In the Matter of)	
)	
Preserving the Open Internet)	GN Docket No. 09-191
)	
Broadband Industry Practices)	WC Docket No. 07-52

Comments of Professor Brett M. Frischmann

Brett M. Frischmann
Associate Professor
Loyola University Chicago
School of Law
25 East Pearson Street
Chicago, IL 60611

April 8, 2010

I strongly support the Commission's efforts to preserve an open Internet. I have published academic papers on the economic and social value of open infrastructure in general and of the Internet in particular. Some of that work has been referenced in the comments filed by others, and I will not reiterate most of the ideas in this brief comment.¹ I have attached the main articles. Instead, I aim to make two brief points:

First, I argue that the FCC must resist falling into the rhetorical trap set by many participants in the debate who attempt to frame the policy debate narrowly in terms of antitrust and regulatory economics. A myopic focus on antitrust and regulatory economics misses other important dimensions at stake in the debate. Essentially, this perspective views the Internet as a mere supply chain of markets. It fails to appreciate that the Internet is a mixed commercial, public, and social infrastructure that supports an incredible variety of market and nonmarket systems and user activities that yield private, public, and social goods. Too many participants in the debate (on both sides) accept the premise that competition would alleviate concerns about discrimination or prioritization by network providers. It would not, as I discuss below.

Second, I offer a particular nondiscrimination rule that differs somewhat from the one articulated by the Commission in the Notice of Proposed Rulemaking. The FCC should prohibit broadband Internet access service providers from discriminating based on the *identity* of the user or use in the handling of packets. Under this approach, user may be defined as sender or receiver; use may be defined as application or content type; handling may be defined as all transport and related services associated with delivery of packets. This simple nondiscrimination rule may seem overly strong in that it appears to rule out a significant range of activities that some might label "reasonable network management." But as I discuss below, this rule is not overly restrictive; rather, it strikes an appropriate balance. It primarily rules out certain fine-grained forms of price or quality discrimination, does not rule out other forms of price discrimination that are not based on user/use identity, such as typical second-degree price discrimination, and does not rule out efficient methods for managing congestion, such as traditional usage-sensitive or congestion pricing. This rule maintains a general-purpose, mixed infrastructure and best preserves the Internet's openness.

¹ I am currently writing a book on the topic: *INFRASTRUCTURE: THE SOCIAL VALUE OF SHARED RESOURCES* (Yale Univ. Press, forthcoming 2011). Prior work that discusses network neutrality includes: Brett Frischmann & Mark Lemley, *Spillovers*, 107 *Columbia Law Review* 257 (2007); Brett Frischmann & Barbara van Schewick, *Network Neutrality and The Economics of an Information Superhighway*, 47 *Jurimetrics* 383 (2007); Brett Frischmann, *An Economic Theory of Infrastructure and Commons Management*, 89 *Minnesota Law Review* 917 (2005). For related work, see Brett Frischmann, *Speech, Spillovers, and the First Amendment*, 2008 *University of Chicago Legal Forum* (2008); Brett Frischmann & Spencer Weber Waller, *Revitalizing Essential Facilities*, 75 *Antitrust Law Journal* 1 (2008); Brett Frischmann, *Cultural Environmentalism and The Wealth of Networks*, 74 *University of Chicago Law Review* 1083 (2007); Brett Frischmann, *Privatization and Commercialization of the Internet Infrastructure: Rethinking Market Intervention into Government and Government Intervention into the Market*, 2 *Columbia Science and Technology Law Review* 1 (2001).

1. *Missing the forest for the trees*

Much of the back and forth among commentators—in comments filed before the Commission but also in the press, at academic conferences, and in the broader public debate—focuses on economics. Framing the debate about sustaining an open Internet in terms of economics has led many people to miss the forest for the trees, and I would like to make sure that the Commission does not fall into that trap.

Numerous comments filed with the FCC and many articles written by economists and attorneys who oppose nondiscrimination rules have a common refrain and rhetorical move. They claim that such rules fly in the face of well-established, widely appreciated, or uncontested lessons in economics. To note just a few examples:

- Opponents of FCC action often use powerful rhetoric to suggest that FCC action would constitute radical intrusion by government into markets that would otherwise be free. This rhetoric reflects a misconception that government and markets as two opposing forces rather than mutually dependent institutions. There is nothing radical about what the FCC has proposed. Similar government regulations exist in other infrastructure industries and have a long history. Claims that deregulation has removed some nondiscrimination rules in some contexts hardly makes the FCC's proposed nondiscrimination rules radical or unprecedented.
- Some would like the Commission and the public to believe that economics provides clear answers to questions about whether regulation is warranted and if so what form such regulation should take. Simply put, economics does *not* provide clear answers. It only provides a limited means for framing the debate and the nature of some of the tradeoffs involved. As Barbara Van Schewick and I explained:

There are many related normative commitments at stake in the network neutrality debate, including market values such as promoting allocative and productive efficiency, innovation, and economic growth but also various nonmarket values such as education and increased participation in cultural and political processes.²

- Many claim that nondiscrimination rules would preclude efficient business practices, such as price discrimination and prioritization. But, as discussed below, the claimed efficiencies of such practices are hardly well-established in general, much less in the particular context of this debate.

Despite the fervor of such claims, the FCC must recognize (1) that these commentators are usually referring to lessons derived from antitrust and regulatory economics,³ and (2) antitrust and regulatory economics provides no clear answers in this

² Frischmann & Van Schewick, at 426-27.

³ Even within antitrust and regulatory economics, the case for nondiscrimination rules is arguably stronger than opponents are willing to admit. The supposedly uncontested lessons are often, in fact, contested.

debate. Reliance on such arguments runs the significant risk of myopia. It is important to understand that antitrust and regulatory economics adopt a partial equilibrium style analysis, focus on competition or the lack thereof within certain markets, assume away incomplete or missing markets associated with public and social goods and thus the spillovers associated with creating and sharing such goods, reduce dynamic considerations about infrastructure evolution resulting from discrimination, and thus paint a woefully incomplete picture.⁴ Through this lens, the Internet is merely a supply chain comprised of markets.⁵

Some argue that modern antitrust and regulatory economics strongly disfavor government intervention into private markets for a variety of well-established reasons. Opponents of network neutrality regulation regularly assert that it is widely understood and well-established within economics that discrimination or prioritization may be efficient and beneficial to consumers; government intervention may stifle incentives to invest in infrastructure; and regulation is likely to be cumbersome, inefficient, wasteful and possibly captured by the regulated entities. These appeals rest heavily on the underlying antitrust and regulatory economic framework and the premises that intervention is only needed when markets are not competitive and that even when markets are not competitive, intervention is only justified in very narrow circumstances where demonstrable harm to consumers in the relevant markets can be shown and not outweighed by efficiency gains.

Competition policy appears to be the fulcrum of the debate in the sense that the debate largely revolves around whether network owners have market power and whether discrimination among data packets causes anticompetitive effects. Proponents of network neutrality regulation claim network owners have market power and the capability and incentives to engage in harmful discrimination. Opponents claim that the relevant markets are competitive and any discrimination practiced by network owners is presumably beneficial. To be fair, participants on both sides of the debate appeal to other considerations.⁶ For example, an important strand of the debate focuses on innovation.⁷

⁴ I briefly address some of these issues below. These problems are addressed more fully in my forthcoming book, *INFRASTRUCTURE: THE SOCIAL VALUE OF SHARED RESOURCES* 60 (Yale Univ. Press, forthcoming 2011).

⁵ The partial equilibrium style analysis involves isolating a particular market (or two) to analyze and holding all others constant and assuming them to be complete and perfectly competitive. The supply chain view of the Internet improperly assumes away an incredible number of incomplete and missing markets, in which users produce and share public and social goods with external effects (spillovers) that are not captured or reflected fully in markets. *Id.*, chapters 3-5.

⁶ A quick survey of the voluminous literature or comments filed in the FCC Proceeding reveals an incredibly wide range of arguments.

⁷ Innovation - commercial and otherwise - is an incredibly important activity furthered by an open Internet. While end-to-end architecture promotes innovation in the first instance, bringing us new, valuable applications and opportunities, it also promotes the continued use of socially valuable but perhaps not commercially valuable applications. See Frischmann, *An Economic Theory of Infrastructure*, at 1015-20; see also Lawrence Lessig, *Re-Marking the Progress in Frischmann*, 89 *Minnesota Law Review* 1031 (2005). On the relationship between Internet architecture and application innovation, see BARBARA VAN SCHEWICK, *INTERNET ARCHITECTURE AND INNOVATION* (MIT Press forthcoming 2010).

Another important strand of the debate focuses on the question of what form of public regulation might be appropriate, antitrust or sector-specific regulation by the Federal Communications Commission. But these strands are heavily influenced by the more basic competition policy framework and its premises.

For example, even the one “intervention” accepted by both sides of the debate—*transparency*—suffers from this myopia and its attendant problems. Transparency is perceived as a solution because it would enable consumers to be effective market participants, capable of disciplining access providers by making a fuss about harmful discrimination and switching providers (when actual competition exists). There at least two problems with this “solution.” First, it presumes consumers have the time, inclination, and capability to process and react to the information made available.⁸ Second, and more important, it presumes complete markets in which consumer demand effectively manifests societal demand. Simply put, transparency alone would not be enough to preserve an open Internet. Rather, it would provide a false sense of security and reinforce the supply chain view of the Internet.

The focus on competition is misguided because it distorts the debate by ignoring demand-side issues—essentially, by assuming that private demand reflected in markets fully reflects social demand. As a result, it fails to appreciate that the social value of the Internet greatly exceeds its market value; that a tremendous amount of the social value derives from activities by and among users associated with the producing, sharing, and reusing ideas and other intellectual and social goods;⁹ and that relying on unfettered market allocation of infrastructure access and use runs the risk of significant demand side market failures.¹⁰ There are two sets of concerns, which I discuss at length in my published work and forthcoming book: First, concerns about undersupply and under-use of infrastructure to produce infrastructure-dependent public and social goods and thus underproduction of those goods. Second, concerns that infrastructure development may be skewed in socially undesirable directions, for example, if private infrastructure owners prematurely optimize infrastructure for uses that they expect will maximize their private returns and in doing so choose a path that forecloses production of various public or social goods that would yield greater net social returns.

It is important to appreciate that Internet use generates many different types of externalities, including third-party effects associated with incomplete or missing

⁸ For further discussion, see, for example, Barbara van Schewick, Written Testimony at the Federal Communications Commission’s Second Public En Banc Hearing on Broadband Network Management Practices at Stanford University, Stanford, CA on April 17, 2008, Docket No. 07-52, p. 4; VAN SCHEWICK, *INTERNET ARCHITECTURE AND INNOVATION*, pp. 261-264.

⁹ Much of the social value derived from the Internet is associated with user participation in activities that generate public and social goods that yield spillovers *offline*. For example, political discussion and ideas exchanged online involve benefits online and offline. In other work, referenced above (note 1), I discuss this at length and give various examples. See also Christiaan Hogendorn, *Spillovers and Network Neutrality* (2010).

¹⁰ See Frischmann, *An Economic Theory of Infrastructure*. For a more detailed analysis of these concerns, see FRISCHMANN, *INFRASTRUCTURE: THE SOCIAL VALUE OF SHARED RESOURCES* ch. 4 & 5.

markets.¹¹ The Internet facilitates productive behaviors by users that affect third parties, including other users, platform providers, and *non-users*—that is, people offline. The third-party effects often are accidental, incidental, and not especially relevant to providers or users. Such effects are not reflected fully in market demand (willingness to pay), but nonetheless may be quite important from a societal perspective.

Critically, even if we assume that we have robust competition in the infrastructure markets, the case for network neutrality remains quite strong. Competition alone does not alleviate the demand-side concerns. It does not ensure an efficient allocation of resources. It does not assure us the Internet environment that maximizes social welfare. Competition does not address these interests for essentially the same reasons that antitrust law is orthogonal to environmental law—antitrust law does not address market failures associated with externalities, whether environmental pollution (negative externalities) or the production, sharing and productive reuse of public and social goods (positive externalities). Indeed, it is “well-established” in economics that competitive markets overproduce pollution and underproduce public and social goods.¹²

Of course, this does not mean that market power and anticompetitive effects do not matter or that recognizing the existence and importance of spillovers associated with user-generated public and social goods provides easy, determinate answers. The network neutrality debate is and must be complicated, but it should not be reduced to a competition policy framework. It needs to grapple with the demand-side issues, including the infrastructural nature of the Internet and the significant role of users in generating social value.

As I explain in considerable detail elsewhere,¹³ the Internet is a mixed commercial, public, and social infrastructure that supports an incredible variety of market and nonmarket systems and user activities that yield private, public, and social goods. The social value of the Internet greatly exceeds its market value precisely because it enables many systems (markets and non-markets) to function and support productive activities by many different types of end-users. Consider what makes the Internet valuable to *society*. It is very difficult to estimate the *full* social value of the Internet because of the wide variety of downstream uses that generate public and social goods.

¹¹ When they do surface, discussions of spillovers can be remarkably reductionist, focusing exclusively on spillovers that involve benefits flowing between networks, content providers, and users—essentially, benefits fully captured online. *See, e.g.*, Declaration of Marius Schwartz, at 24-25, filed by AT&T (Jan 14, 2010); Frischmann & Van Schewick, at 399-400, 402-03, 427 (expanding on this point).

¹² Conventional economic solutions to the underproduction of public goods, such as directing subsidies to public goods producers, do not work very well in this context because of the incredible variety of producers and of public goods, and the predictable failure of government in choosing how to direct subsidies in this environment. *See* FRISCHMANN, *INFRASTRUCTURE* ch. 5. Further, the solutions may not work as well for social goods that are jointly produced in a more distributed fashion. *See, e.g.*, YOCHAI BENKLER, *THE WEALTH OF NETWORKS: HOW SOCIAL PRODUCTION TRANSFORMS MARKETS AND FREEDOM* (2006); FRISCHMANN, *INFRASTRUCTURE* ch. 3 (discussing four types of social goods—nonmarket goods, merit goods, social capital, and irreducibly social goods).

¹³ For the sake of brevity, I refer the Commission to the publications cited in the first note.

The Internet is integral to the lives, affairs and relationships of individuals, companies, universities, organizations, and governments worldwide. It is having significant effects on fundamental social processes and resource systems that generate value for society. Commerce, community, culture, education, government, health, politics, and science are all information- and communications-intensive systems that the Internet is transforming. The transformation is taking place at the ends, where people are empowered to participate and are engaged in socially valuable, productive activities. The Internet opens the door widely for users. End-users actively engage in innovation and creation; conduct commercial and non-commercial exchange; speak about anything and everything; maintain family connections and friendships; debate, comment, and engage in political and non-political discourse; meet new people; search, research, learn, and educate; and build and sustain communities. These are just some of the productive activities that generate substantial social value, value that often evades observation or consideration within conventional economic transactions. Participation in such activities results in external benefits that accrue to society as a whole, both online and offline, that are not captured or necessarily even appreciated by the participants.

The Commission should recognize that *activities on the Internet* always involve *interactions* among *end-users*. The interactions may be commercial, educational, social, political, and so on, and the end-users may be individuals, corporations, government actors, or other entities. This is another place where the antitrust/competition policy framework distorts the debate. The framework leads to misconceptions of the players involved. Specifically, the debate focuses on how network neutrality would affect access providers, content providers, and consumers. *Who are the content providers?* Google, of course, and other large companies like Google; big companies that deliver a lot of content to consumers. But this framing conflates too much, by understating the role of consumers as producers and also by creating a false distinction between Google and other end-users. Consumers are users that produce applications and content, and Google is a consumer and user and producer, like me. Of course, we are different in some ways, but not necessarily in a way that matters. (I could be the next Google (I wish!) or the Google of my own “space” on the net.) The key point is that viewing the Internet as a means for distributing content, applications and services to consumers biases the debate in a way that, again, misses the forest for the trees. The Internet is much more than a low cost delivery system for content providers or a two-sided market with networks acting as intermediaries between content providers and consumers. It is a mixed commercial, public, and social infrastructure that supports an incredible variety of market and nonmarket systems and user activities that yield private, public, and social goods.

2. *A proposed nondiscrimination rule*

The Commission is rightly concerned about sustaining an open Internet. A targeted nondiscrimination rule is an appropriate intervention because it would both preclude differentially allocating and prioritizing access and use of the Internet on the basis of expected private returns and also limit infrastructure evolution or optimization on

that basis.¹⁴ In this section, I discuss my proposed nondiscrimination rule, which precludes discriminating based on the *identity* of the user or use.

The rule proposed in FCC NPRM invokes six principles, all of which are modified by exception for “reasonable network management.” The basic principles make sense, but the (undefined) exception swallows the rule. Many other comments have addressed this issue.¹⁵ Rather than rework those discussions, I propose an alternative formulation.

The FCC should consider prohibiting broadband Internet access service providers¹⁶ from discriminating based on the *identity* of the user or use in the handling of packets. User may be defined as sender or receiver; use may be defined as application or content type; handling may be defined as all transport and related services associated with delivery of packets. This simple formulation of a nondiscrimination rule may seem overly strong in that it appears to rule out a significant range of activities that some might label “reasonable network management.” Of course, it all depends on what that label applies to. As I discuss below, the rule primarily rules out certain fine-grained forms of price or quality discrimination, does not rule out other forms of price discrimination that are not based on user/use identity, such as typical second-degree price discrimination, and does not rule out efficient methods for managing congestion, such as traditional usage-sensitive pricing or congestion pricing. After discussing how the rule would impact pricing schemes, I explain why disabling application-based prioritization is not as horrible as opponents suggest.

In essence, the proposed rule precludes many forms of value-based price discrimination.¹⁷ It does not preclude various forms of cost-based differential pricing, including variable load pricing, congestion pricing, and usage-sensitive pricing. The reason why is that these and other forms of cost-based differential pricing do not discriminate on the basis of the identity of users or their specific activities (uses) and instead discriminate based on quantity of infrastructure use, capacity utilized, and the marginal cost of such use, taking into account contextual details such as timing and available system capacity but not identity characteristics. Put simply, to implement most forms of cost-based differential pricing, an infrastructure owner need not know *who is doing what*. Rather, the focus is on *when* and *how much*. In sharp contrast, price discrimination relies directly on identity characteristics that aim to best approximate individual users’ subjective valuation of infrastructure use—*who and what are essential to the discrimination or prioritization scheme*.

¹⁴ I discuss this extensively in prior publications and chapter five of my forthcoming book. For the sake of brevity, I refer the Commission to the publications cited in the first note.

¹⁵ See, e.g., Comments of Sony Electronics, at p. 9; Comments of Free Press, at pp. 84-85, 91; Comments of Public Interest Commenters, p. 35; Comments of Texas Public Utility Counsel, at p.5; Comments of Google, at 73.

¹⁶ Some have suggested that the Commission should not limit its rule to access providers. See Comments, MSU IP & Communications Law Program.

¹⁷ For ease of discussion, I refer to price discrimination and leave aside other forms of discrimination (e.g., quality).

Value-based discrimination is at the core of the rule because it is estimated appropriable value that drives private allocation, prioritization and optimization decisions and potentially leads the infrastructure to evolve in a manner that forecloses production of various public and social goods and thereby reduces social welfare. The proposed nondiscrimination rule precludes differentially allocating and prioritizing Internet access and use on the basis of appropriable value. This can be a more effective—albeit blunt—means for supporting the production of public and social goods than targeted subsidies.¹⁸ The rule also has important dynamic implications. It maintains flexibility and the generic nature of the Internet. The proposed nondiscrimination rule effectively functions as a social option, which makes economic sense because of persistent and systematic uncertainty about the future sources of both market and social value.¹⁹ It precludes premature optimization by infrastructure owners, preserves the “evolvability” of the Internet,²⁰ and thereby supports experimentation by users, increases the range of potential value-creating activities, leaves room for unforeseen innovations, markets and value-creating activities to emerge, and facilitates learning over time.²¹

Of course, disabling value-based discrimination is not necessarily costless. Conventional economic analysis of price discrimination suggests that in some cases, price discrimination can improve aggregate social welfare (relative to uniform pricing) by increasing output and, by virtue of the increased private returns, improving incentives to invest in the infrastructure. Yet, like many conventional analyses of infrastructure markets, the conventional economic analysis of price discrimination adopts a partial equilibrium style analysis, assumes away incomplete or missing markets associated with dependent public and social goods, reduces dynamic considerations about infrastructure evolution resulting from discrimination, and thus paints an incomplete picture. I do not intend to challenge the substantial general economics literature on price discrimination here, however. Instead, I make a few general observations that are applicable to this context.

¹⁸ Although the conventional economic solution to supporting production of a public good is targeted subsidies, this solution does not work well in this context. See *supra* note 12. Disabling price discrimination at the infrastructure level effectively creates cross-subsidies among uses, which can be a more efficient means of supporting a wide variety of public and social goods. I develop this point in a series of articles. See sources cited in note 1. Other scholars have picked up on the idea. See Robin S. Lee and Tim Wu, *Subsidizing Creativity Through Network Design: Zero Pricing and Net Neutrality*, *Journal of Economic Perspectives*, Vol. 23, No. 3, pp. 61-76, Summer 2009 (making this argument in the context of the network neutrality debate). For extensive discussion, see FRISCHMANN, *INFRASTRUCTURE* chs. 3, 5.

¹⁹ See *id.* See also Mark Gaynor, *Network Services Investment Guide: Maximizing ROI in Uncertain Times* (2003); Mark Gaynor & Scott Bradner, *Statistical Framework to Value Network Neutrality*, 17 *MEDIA L. & POL'Y* 24 (2007); VAN SCHEWICK, *INTERNET ARCHITECTURE AND INNOVATION*.

²⁰ See *id.*.

²¹ See FRISCHMANN, *INFRASTRUCTURE* ch. 5.

Truly perfect price discrimination would eliminate many allocation concerns.²² Despite the superficial appeal of perfect price discrimination, there are (at least) three reasons to be wary of it.

- First, perfect price discrimination is *not* feasible in the real world. It simply requires an inordinate amount of fine-grained information about consumers. The unattainable but theoretically appealing ideal of perfect price discrimination operates more as a distracting red herring than a useful analytical baseline.
- Second, the path to perfect price discrimination may be fraught with peril for society.²³ Pricing practices in many industries evolve over time.²⁴ When unchecked by competition or regulation, producers often drive toward price discrimination. To progress on this path, producers often depend upon adapting law, technology and other institutions to their needs. For example, to facilitate price discrimination, producers may lobby for changes in the law (for example, relaxing antitrust scrutiny of patent licensing), invest in technologies (for example, deep packet inspection), or alter conventional relationships with consumers (for example, by collecting massive amounts of personally identifiable information). These steps may involve considerable social costs and unintended consequences, such as a dramatic diminution in privacy to satisfy the immense information requirements for price discrimination.
- Finally, perfect price discrimination would eliminate cross-subsidies that arise when infrastructure is managed in a nondiscriminatory manner. This may have dynamic effects on production of public and social goods in situations where the consumer surplus that is converted to producer surplus would otherwise be necessary to support certain productive activities by infrastructure users. This latter point is completely missing in the literature because infrastructure-dependent markets are generally assumed to be complete.²⁵

Given the unattainable nature of perfect price discrimination, the appropriate scheme to consider is imperfect price discrimination. Imperfect price discrimination, which includes second-degree and third-degree price discrimination, is quite common in

²² Perfect price discrimination, also known as first-degree price discrimination, means that each unit of output is sold to the person that values it most and at that person's maximum willingness to pay. Thus, price would vary both by person and by unit of output, being perfectly calibrated to each individual's preferences. The producer would capture the entire consumer surplus. Absent congestion, all users that desired access to the infrastructure would be given access at their respective willingness to pay and thus deadweight losses would be eliminated. To be clear, perfect price discrimination would involve no deadweight losses because no one would be priced out: Marginal consumers would only have to pay what they would be willing to pay, constrained only by the marginal costs.

²³ For further discussion, see Frischmann, *An Economic Theory of Infrastructure*, at 978-80. For a more detailed analysis of these concerns, see FRISCHMANN, *INFRASTRUCTURE* ch. 5.

²⁴ See Andrew Odlyzko, *The Evolution of Price Discrimination in Transportation and Its Implications for the Internet*, 3 REV. NETWORK ECON. 323 (2004).

²⁵ See Frischmann, *An Economic Theory of Infrastructure*, at 975-78, 1020-21; *Spillovers*, at 296-98. For a more detailed analysis, see FRISCHMANN, *INFRASTRUCTURE* ch. 5.

various industries, including many infrastructure industries. Second-degree price discrimination occurs when price per unit varies based on the amount purchased; for example, when bulk purchasers receive a discount. Third-degree price discrimination occurs when price per unit varies based on consumer groups; for example, when students or seniors receive a discount. Economists generally recognize that the welfare implications of imperfect price discrimination are ambiguous and vary considerably by context.²⁶

Critically, my proposed rule does not preclude all forms of price discrimination. It precludes discriminating on the basis of the identity of the infrastructure user or use. Second-degree price discrimination, for example, may or may not run afoul of this nondiscrimination rule, depending on how the scheme operates. Second-degree price discrimination may present all infrastructure users with the same price schedule for the same basic service with price variations based on the quantity of use or capacity consumed; users decide what to choose from the menu based on their anticipated demand. The price schedule may differentiate among customer classes, based on differences in the quantity of infrastructure capacity demanded, the costs of providing service, and group elasticity of demand, meaning that those charged a higher per unit cost are less likely to alter their consumption. Second-degree price discrimination of this sort is very common in infrastructure industries. When prices vary nonlinearly with the amount purchased or capacity consumed—for example, when high volume customers pay less on a per unit basis than low volume customers for identical service—the discrimination is not identity-based. *Who is doing what* is not really important to the categories or the pricing scheme; the focus is on *how much capacity* and *at what cost*. In general, this does not conflict with the proposed rule.

Third-degree price discrimination, by contrast, tends to categorize consumers based on their identity, specifically based on identity characteristics that serve as effective proxies for consumers' subjective valuations. The categories are not based on cost. *Who is doing what* is central to the scheme. Third-degree price discrimination generally conflicts with the proposed rule. Of course, so do more fine-grained identity-based price discrimination schemes.

As noted, disabling identity-based price discrimination is not necessarily costless. It may involve a tradeoff. The two principal potential advantages of price discrimination are (i) increased output and thus reduced deadweight losses when compared with uniform pricing, and (ii) increased profits for infrastructure providers that may improve incentives to invest in the supply, maintenance, and improvement of infrastructure. Neither of these potential advantages appears to be nearly as significant as claimed by many in the debate.

²⁶ See JEAN TIROLE, *THE THEORY OF INDUSTRIAL ORGANIZATION* 139, 149 (1988) (concluding that the welfare effects of imperfect price discrimination—technically, second and third degree price discrimination—are “ambiguous” and may be “socially suboptimal,” depending upon the context); DENNIS W. CARLTON & JEFFREY M. PERLOFF, *MODERN INDUSTRIAL ORGANIZATION* 448–50 (1990).

First, the claim that output would increase if network providers can prioritize traffic is somewhat misleading. The case for prioritization depends on an assumption of congestion. In the absence of congestion, there is no need to prioritize traffic (other than rent extraction).²⁷ If the network is congested and infrastructure capacity is scarce, then output cannot be increased without making congestion worse or imposing congestion costs on others. If A and B are two use(r)s vying for access and use of a congested resource, prioritization of use(r) A requires deprioritization of use(r) B. The argument for prioritization is really an argument to allocate scarce capacity to uses that would suffer greater consequences from congestion—whether attributed to delay or jitter or packet loss—and thus would realize greater benefits from prioritization. That is, it is not really an argument for increasing output measured by the number of users, uses, or even packets delivered during periods of congestion; prioritization does not eliminate or minimize congestion.²⁸ In fact, as various comments have pointed out, prioritization creates perverse incentives for network providers to sustain congestion and underinvest in capacity expansion.

Second, the oft-stated claim that increased revenues from price discrimination are necessary to investment incentives is speculative and self-interested. Major providers, such as AT&T and Verizon, have invested considerable sums without any assurance that they would be capable of prioritizing traffic to obtain increased revenues. Of course, they would love to capture more of the surplus, but there is no reason to think they are entitled to it, and it is hardly proven that they need it to justify infrastructure investment.²⁹ I have addressed this argument elsewhere, and many other commentators have addressed this point.³⁰

²⁷ I do not discuss identity-based price discrimination schemes that are divorced from prioritization during congestion and are essentially rent extraction schemes aimed at capturing a greater portion of the surplus derived from various Internet activities. For an analysis of these schemes, see Barbara van Schewick, *Internet Architecture and Innovation* (MIT Press forthcoming 2010), pp. 273-278.

²⁸ Users A and B may adjust their consumption patterns based on how congestion costs are distributed (shared equally or pushed to B because B is less sensitive to timing) and such adjustments may or may not lead to increased output over the long run if there is sufficient off peak capacity.

²⁹ Moreover, “if in fact incentives to invest in infrastructure are suboptimal and need to be improved, there are alternative ways of doing so that do not entail discrimination on the basis of content or application. Some viable options include direct subsidization of infrastructure expansion, tax incentives to support the same, cooperative research and development projects, and joint ventures. In addition to improving incentives for private provision of infrastructure, government provision of infrastructure is another viable option for the last mile. In fact, municipal provision of broadband access to the Internet has gained significant momentum in recent years. But we are skeptical that any of this is necessary. Telecom companies radically overbuilt infrastructure in the late 1990s without any of the rights they now demand, and a number of companies are competing to build free public wireless networks in major cities throughout the country, again without any such guarantee.” Frischmann & Lemley, *Spillovers*, at 297 n.147.

³⁰ See, e.g., Comments of PAETEC, at 21-25; Comments of Google, at 37-41, Appendix A; Comments of Open Internet Coalition, at 30-33.

Let me close with two additional points. First, investment in capacity expansion and pricing sensitive to usage and congestion³¹ can go a long way in resolving many of the perceived congestion problems on the Internet.³² End-to-end design works reasonably well because all packets in fact are the same from a delivery cost and congestion cost perspective. Of course, all packets are not the same from the user valuation perspective. Users derive value from the higher layer uses of delivered packets, and different uses are more or less sensitive to the consequences of congestion. As discussed above, the economic case for prioritization depends on the existence and persistence of congestion. But there are good reasons to direct policy toward eliminating or at least minimizing congestion, primarily through expanding capacity and implementing usage-sensitive or congestion pricing, rather than accepting (and possibly encouraging) persistent congestion. This means that solving congestion problems through means that do not discriminate based on the identity of the user or use may be sufficient and desirable because it pushes providers to focus on capacity expansion and usage-sensitive pricing rather than prioritization.

Second, my proposed rule certainly could admit exceptions for categorical discrimination against traffic that is harmful to the network itself *if such harm can be shown*. However, exceptions based on congestion should not be allowed. Further, exceptions based on harm to users because of the content itself (e.g., spam) may not be needed. It remains unproven (at least, to my knowledge) that such issues cannot be dealt with effectively at the ends. Innovation at the ends has proven remarkable at addressing problems of all sorts. As Mark Lemley and I noted, “even with a dumb architecture, innovators have figured out how to provide certain degrees of quality of service at the periphery of the network and how to make certain latency-sensitive applications, such as IP telephony, work.”³³ Further, it is worth noting that in many situations, access networks operate as end systems, particularly where access networks run server-based applications and provide related services for customers. For example, a mail server is an end point, whether owned and managed by an access network such as Comcast or a customer of the access network. Blocking or filtering traffic at the mail server based on the originating IP address is a routine method for dealing with spam that is end-to-end compliant—it involves one end dealing with another end, and the discrimination does not occur at the transport layer. This does not conflict with the proposed rule. Still, if an access network blocks or otherwise discriminates in the treatment of traffic not destined for its own servers, such as peer-to-peer traffic destined for a customer’s home computer, then the access network would violate end-to-end and the proposed rule.

³¹ For a detailed discussion, see Brett Frischmann & Barbara van Schewick, *Network Neutrality and The Economics of an Information Superhighway*, 47 *Jurimetrics* 383, 392-409 (2007).

³² Of course, congestion will not disappear altogether, and congestion pricing (much like quality of service assurances) likely cannot be implemented system-wide because of the many different networks and resources involved, although usage-based pricing can and probably should be implemented system-wide.

³³ See Frischmann & Lemley, *Spillovers*, at 295 n.143.