COLLEGE OF LAW

University of Nebraska Lincoln

Justin (Gus) Hurwitz
Assistant Professor of Law
(402) 472-1255
ghurwitz@unl.edu

Before the
Federal Communications Commission
Washington, D.C. 20554

In the Matters of	)	
	)	
Protecting and Promoting the Open Internet	)	GN Docket No. 14-28
	)	
Framework for Broadband Internet Service	)	GN Docket No. 10-127
	)	

Comments of Justin (Gus) Hurwitz, Assistant Professor of Law, University of Nebraska College of Law[1]

These comments are respectfully submitted in response to the Commission's May 15, 2014, Notice of Proposed Rulemaking in GN Docket Number 14-28, relating to the Open Internet proceedings, and related matters. Also included with these comments are copies of recent short writings relating to the issues presented in the Notice.

To summarize, these comments argue that the Commission should not adopt comprehensive "Open Internet" rules, but rather should promulgate general guidelines to provide notice to industry participants (be they consumers, last-mile or backbone carriers, or edge providers) about the general classes of conduct of concern to the Commission. Should subsequent conduct raise concerns under these guidelines, the Commission should take strong and swift enforcement action, proceeding on a case-by-case basis under Section 706.

This approach, which is captured in the proposed "commercially reasonable" standard, is grounded in modern principles of administrative law and procedure; is sufficient to protect consumers and police potentially problematic behavior; is pragmatic both in that it is less

[1] The author teaches and researches in areas relating to telecommunications law and policy, Internet- and cyber-law, and the regulation of technology; has previously worked in the Federal government in a capacity relevant to the present NPRM; has degrees in both law and economics and has researched and published in technical fields relating to modern telecommunications technology. Further information about the author is at http://law.unl.edu/facstaff/faculty/resident/ghurwitz.

P.O. 830902 / Lincoln, NE 68583-0902
(402) 472-2161 / FAX (402) 472-5185

likely to result in further years of Open Internet rulemaking efforts (and concomitant uncertainty) than other alternatives being considered by the Commission and in that it is more likely to survive judicial challenge than those alternatives. It allows for the continued development of new pro-consumer businesses and business models on the Internet while allowing the Commission to investigate and take action against conduct that may harm consumers – in this sense it is unlike other alternatives under consideration that could foreclose the development of pro-consumer businesses and business models in the interest of preventing hypothetical (and often nonsensical) consumer harm. This approach maintains flexibility that is essential to avoid manipulation, capture, and arbitrage by firms that would use the Commission's rules to profit, possibly at the expense of consumers.

These comments proceed in four parts. The first part discusses the paramount importance of consumer welfare in whatever rules the Commission adopts. In addition to common points about the importance of consumer welfare, this discussion argues that consumer welfare urges that the Commission adopt rules that will not result in immediate judicial challenge and additional years of development and litigation, that it requires the FCC to play an affirmative role in educating consumers, and that it requires the development of a dispute and complaint resolution mechanisms that is timely, accessible, and responsive to consumers.

The second part briefly surveys the economics of vertical integration, innovation, and multi-sided markets, to argue that legal standards are preferable to rules. The literature on these topics consistently demonstrate that vertical integration, platform "openness," and multi-sided business models can have inconsistent and unpredictable effects on consumer welfare. Any rule, therefore, that prophylactically mandates or forecloses the adoption of broad classes of business structures or conduct is, almost by definition, arbitrary, and its adoption capricious.

The third part discusses the Commission's authority to implement Open Internet rules. It makes the basic point that adoption of Title II as a legal basis would almost certainly result in additional years of industry and consumer uncertainty as the Commission's rules are challenged by interested parties in court and narrowed through forbearance. This substantial cost comes with little gain, as Title II does not offer substantially more legal authority than Section 706, which is already sufficient to the Commission's legal goals.

The third part's discussion also considers the question of implementation – a focus that has thus far largely been absent from discussion of the Commission's rules. It argues that *ex post* case-by-case enforcement of standards is preferable to the development and enforcement of strong *ex ante* rules, and explains the legal requirements to proceed along such a path.

The final part offers a technical discussion of statistical multiplexing and its relationship to the Notice's consideration of minimum levels of access and nondiscrimination. The level of technical detail is modest, but makes a critical point with respect to concerns about so-called "fast lanes": in statistically multiplexed systems (such as the Internet), when one user (or firm) pays another for committed capacity, all other users of the network benefit when that user's capacity is not being used. Conversely, when one user consumes a statistically disproportionate amount of capacity, all other users on the network are harmed. Understanding the technical principles that undergird how Internet resources are shared is essential to understanding concepts like "best effort," minimum levels of access, and harmful discrimination.

I. First, do no (consumer) harm

It is common – rote bordering on trite – to begin comments such as these by asserting the primacy of the consumer in the regulatory calculus. But it is important to do so. Consumer welfare is the lodestone that should guide all government policy; and, in the form of the "public interest, convenience, and necessity," it is expressly the Commission's guiding principle.

Consumer welfare is important to the Commission's proposed rules in several ways. Foremost, regulation of business models and other conduct online has direct and indirect effects on consumer welfare. Good rules will, on net, increase consumer welfare; bad rules will, on net, reduce it. As is argued later, the economics literature consistently demonstrates that the consumer welfare effects of the sort of business models and conduct addressed by the proposed rules are uncertain – which means that any firm rules adopted today have the potential to harm consumers.

But consumer welfare is also implicated in other ways by the proposed rules – as well as by the broader Open Internet/Network Neutrality debate. Three ways the NPRM affects consumer welfare are discussed here: consumers getting caught in fights between content and distribution; consumer uncertainty and confidence in the industry; consumer access to and understanding of regulatory protections. Each of these concerns should be incorporated into the Commission's analysis and ultimate adoption of any rules.

Consumers often get caught in the middle of fights between content owners and content distributors. We see this today in the fights between Netflix and ISPs, and in retransmission consent disputes. We have also seen it historically, for instance in disputes between long distance and local exchange carriers in the early 1900s, and (more poignantly) between Western Union and the Associated Press in the late 1800s.

Fights between content owners and distributors are often over how to divvy up rents extracted from consumers (i.e., profits). Where this is the case, the actual effect of the dispute on consumers is relatively small – though the resources expended in the fights are deadweight losses, the fight itself is over the distribution of existing rents, not the creation of new rents. This changes, however, when firms involve consumers in these disputes, either by withholding access to consumers or content (e.g., distributor refusing to carry or degrading content or a content provider refusing to provide it), or by involving consumers as parties to the dispute (e.g., by encouraging consumers to petition the government for regulation, based on incomplete or biased information). Importantly, both consumers and government naturally focus their initial scorn on content distributors – perceived as mere middlemen skimming undeserved profit from the valued and desirable wares of the content providers; but it is often later discovered that the content owners are the actual malefactors in the dispute with the distributors.[2]

This demonstrates two important points. First, firms often use the regulatory process for their private gain – and this is often done at the expense of consumers. For this reason, antitrust enforcers are generally reluctant to credit complaints by a firm's competitors, or those who could otherwise gain competitive advantage through regulatory action. And, second,

---

2 *See, e.g.,* Gus Hurwitz, *Should we regulate firms we just don't like?*, TechPolicyDaily.com (November 6, 2013) (citing and discussing Menahem Blondheim, *Rehearsal for Media Regulation: Congress Versus the Telegraph-News Monopoly, 1866-1900*, 56 Federal Communications Law Journal 299 (2004)).

consumer effects are often difficult to assess and run contrary to expectations. Again drawing from the antitrust experience, this is the basic reason that most antitrust analysis has moved from *per se* rules, proscribing certain conduct as permissible or impermissible, to Rule of Reason analysis, under which antitrust claims are examined on a fact-specific, case-by-case basis.

These concerns are particularly important in media-related industries, where consumer sentiment may be naturally biased against distributors and in favor of content owners. This suggests in part that the Commission should be cautious in accepting consumer expressions of concern about maintaining the Open Internet, especially where they assert concerns about distributor relationships with content owners as the basis for their concerns.

This also suggests that the Commission incorporate consumer concern into its welfare analysis. The Commission leadership and staff know as well as anyone that the rules at issue in the NPRM have captured substantial consumer attention. Ongoing fights about network neutrality create substantial uncertainty and concern for consumers. No matter whether these concerns, or consumers' suggested approaches to resolving them, are reasonable, they can legitimately be viewed as consumer harm.

There are two things that the Commission should do in response to these concerns. First, in adopting its Open Internet rules, the Commission should consider how the implementation of those rules will affect consumer understanding of and confidence in the industry. For instance, between otherwise equal options, the Commission should favor rules that are more straightforward or timely to implement, that would reduce firms' ability to strand or otherwise strategically use consumers, and that are less likely to immediately end up back in court. This last point is particularly important given the Commission's consideration of using Title II as a basis for its authority. Should it elect to do so, it seems likely that the new rules will quickly be challenged, resulting in several months or years of uncertainty before the rules can be implemented. Even absent such litigation, the Commission will likely face years of discussion and litigation over the forbearance of various Title II requirements. This outcome would constitute substantial consumer harm.

Second, in developing and implementing its rules, the Commission must take seriously the need for consumers to understand the rules, report their concerns, and be provided with explanatory feedback about these concerns. The NPRM considers the creation of an ombudsperson; this proposal should be firmly embraced. A central task of this office should be to educate consumers – particularly by providing substantive explanatory feedback when action is not taken in response to a complaint. As part of this task, information and statistics about complaints and responses to them should be compiled and maintained publicly. It is also important that this office be non-political.

II. Economic and technical complexity urge use of standards over rules

A basic assumption of the NPRM is that it is possible to craft pro-consumer Open Internet rules. While there may be some business practices and other forms of conduct that are sufficiently harmful (or beneficial) to consumers to merit *ex ante* treatment – such as conduct that would be clearly problematic under the antitrust laws – the effects of the wider range of practices or conduct subject to the NPRM is generally ambiguous. As such, the Commission would be ill-advised to adopt strong rules – and such rules would necessarily be arbitrary and capricious. Rather, where the Commission feels that the potential for harmful practices or conduct is great, it should provide guidance as to the standards by which it will assess whether specific instances are, in fact, harmful. Such guidance is likely a necessary requirement to satisfy Constitutional Due Process and fair notice requirements, particularly where the Commission may want to seek fines or damages against a firm.

Three types of organizational structures demonstrate this ambiguity: "open" innovation platforms, vertical integration, and multi-sided markets.

Net neutrality advocates, the strongest supporters of Open Internet regulation, often assert that the Internet has thrived because it is "open." Indeed, this idea is captured by (or perhaps has captured) the Commission's own caption for this docket: Protecting and *Preserving the Open Internet*. The reality, however, is that "openness" is neither necessary nor sufficient for the sort of growth that the Internet has seen and fostered – indeed, it can limit such growth. Moreover, the Internet never has been open in the way that advocates suggest.

II.a The Internet as an "open" platform

On the first point, literature and experience amply demonstrate that "open" platforms, or general purpose technologies generally, *can* promote growth and increase social welfare; they *can* also create limits.[3] These limits originate along myriad vectors. Developing open systems can impose costs on initial developers, both in their development of the system and their education of third parties in how to interface with it. It can ossify a system, as subsequent changes will have negative spillovers for third parties. It can fragment a system, if subsequent changes are made by the initial developer or subsequent developers choose to fork a project. It can impose technical costs, as an open system may need to be "more robustly" (e.g., inefficiently) engineered in order to support open interfaces. In particular, to be truly "open," a system may need to make internal variables and functions available externally (to those outside the system), even where they are best kept internal to the system[4] – this can lead to substantively inefficient system design, or inefficient use of development resources. An example of this can be seen in the architectural differences between the Internet Protocol stack, which is a four layer model, and the OSI reference model, which is a seven layer model.[5] The OSI model is unquestionable more "open," in that it allows a greatly larger number of interface pairings and the export of a greater number of internal variables and functions to facilitate those pairings; the OSI model, however, is rarely used, largely because it is grossly over-engineered. A protocol designed on the OSI model would cost more to develop, be harder to maintain, and have worse technical performance than one designed on a more streamlined (and less "open") model.[6]

---

3 The seminal cite is Bresnahan & Trajtenberg , *General Purpose Technologies "Engines of Growth?"*, 65 J. ECONOMETRICS 83, 94–96 (1995). See also Christopher Yoo, *Modularity Theory and Internet Policy*, at 33 (2013) (discussing same), *available at* http://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=1467& context=faculty_scholarship.

4 See Yoo, at 22–24 (discussing interdependency and information hiding between modules).

5 The two models' layer approaches represent somewhat different purposes, but are broadly comparable, if on general terms.

6 As a simple, but important example: the most computationally time-consuming process performed in the TCP/IP stack is the passing the payload (the actual data being sent between endpoints) between layers. In a truly "open" stack, IP requires data to be copied three times across the memory bus; for large packets, this can dramatically increase latency. An integrated, "closed," stack, on the other hand, can perform this same function with a single copy operation. Many high-performance IP stacks do precisely this, often integrating the IP or even TCP processing functions with the network interface ("offloading" the IP or TCP processing). Importantly, this involves a form of deep packet inspection, where the middle layers of the IP stack need to both look down the stack into a packet's

Beyond the academic literature, much of which is grounded on theory or models, the relative merits of open versus closed systems has been vividly played out in practice. The classic fights here have played out between, e.g., IBM (a closed system model) and IBM-compatible PCs (open systems); between Apple (a notoriously closed system) and other OS models (e.g., Linux, an open system; and Microsoft, an intermediary model); between application vendors (e.g., traditional closed-source models and the open-source development model); and between different mobile platforms (e.g., between Apple's closed iOS and Google's partially-open Android platform). The results of these fights are generally well known, and demonstrate the indeterminacy of the value of "openness." Apple is the clear example demonstrating the success of closed models. The Apache web server is a leading example of the value of openness.[7] Google's Android is an enlightening example. Google develops the core OS on an open source basis, but retains closer control over the core suite of applications that run on the platform. One should also consider Wikipedia and its struggles to maintain its quality as an open platform over time.[8]

The conclusion to draw from this discussion is that, while open platforms can generate consumer benefits, there is no *ex ante* reason to believe that on particular platform will do so merely because it is open. To the contrary, closed platforms can prove more valuable than open ones.

The second point to make about the "preservation" of the open Internet is that the Internet simply is not, and never has been, an open platform in the sense that net neutrality advocates take it to be. The TCP/IP stack incorporates many design decisions and compromises that were made that either expressly or incidentally improve or degrade the performance of different types of applications or different users based upon how they connect to the network. Moreover, many of these decisions were made based upon the characteristics of then-available computer hardware – they would likely have been made differently and have had different consequences given today's technology and uses of the network.

This fundamental point has been made clearly and strongly by several research scientists that were intimately involved with the development of the Internet. This is best seen with David Clark. Quoting from a 2009 Communications Daily interview with Clark, Hazlett and Wright recount Clark's description of the early Internet:

> "The network is not neutral and never has been," Clark said, dismissing as "happy little bunny rabbit dreams" the assumptions of net neutrality supporters that there was once a "Garden of Eden" for the Internet. NSFnet, an early part of the Internet backbone, gave priority to interactive traffic, he said: "You've got to discriminate between good blocking and bad blocking."[9]

---

headers and up the stack into system state information, in order to directly place a packet's payload directly into application memory in its initial copy operation.

7 For important recent work on this topic, see Shane Greenstein and Frank Nagle, *Digital Dark Matter and the Economic Contribution of Apache* (2013) ("We argue that these findings point to a large potential undercounting of the rate or return from IT spillovers from the invention of the Internet, and to a large potential undercounting of 'digital dark matter' in general."), *available at* http://www.nber.org/papers/w19507.

8 For discussion of the open platforms' questionable reliance on ongoing input and support from third parties, *see* Smith, et al, *Experiences Enhancing Open Source Security in the POSSE Project*, at 4-5 (2003), *available at* http://users.ics.forth.gr/~sotiris/publications/bookchapters/2posse2003.pdf.

9 Hazlett & Wright, *The Law and Economics of Network Neutrality* (2011) (quoting 2009 Communications Daily discussion with David Clark).

Similarly, Jon Crowcroft explains "the basic realities of the net, which has never been a level playing field for many accidental and some deliberate reasons," concluding that "we never had network neutrality in the past, and I do not believe we should engineer for it in the future either."[10] Others have made this same point from the early years of the FCC's involvement in Network Neutrality.[11]

Much of the discussion about the open Internet – including the proposed rules – focus on openness qua neutral treatment of users, applications, and data by the network. As this is a more constrained understanding of openness than much of the innovation literature cited above, it is useful to note – as Clark, Crowcroft, and others also do – that discrimination in the handling of packets has long been discussed as desirable or necessary in development of the technical standards under which the Internet operates.[12] Similarly, research demonstrates that so-called "non-neutral" treatment can be affirmatively desirable.[13]

II.b. Vertical integration

The literature on vertical integration is related to the literature on innovation and openness – and many of the concerns relating to the NPRM can be expressed in terms of vertical integration. For instance, much of the NPRM is concerned with how a vertically-integrated content & distribution firm treats its own traffic compared to that of its competitors. These are unquestionably valid concerns – indeed, the widest range of potentially problematic

---

10 Jon Crowcroft, *Net Neutrality: The Technical Side of the Debate*, 1 Int'l J. Comm. (2007).

11 Douglas A. Hass, *The Never-Was-Neutral Net and Why Informed End Users Can End the Net Neutrality Debates*, 22 Berkeley Tech. L.J. 1565 (2007).

12 *See, e.g.,* RFC 2475 ("Service differentiation is desired to accommodate heterogeneous application requirements and user expectations, and to permit differentiated pricing of Internet service."); RFC 2638 (discussing paid prioritization, saying: "It is expected that premium traffic would be allocated a small percentage of the total network capacity, but that it would be priced much higher."); RFC 1633 ("real-time applications often do not work well across the Internet because of variable queuing delays and congestion losses. The Internet, as originally conceived, offers only a very simple quality of service (QoS), point-to-point best-effort data delivery. Before real-time applications such as remote video, multimedia conferencing, visualization, and virtual reality can be broadly used, the Internet infrastructure must be modified to support real-time QoS, which provides some control over end-to-end packet delays." ... "The first assumption is that resources (e.g., bandwidth) must be explicitly managed in order to meet application requirements. … An alternative approach, which we reject, is to attempt to support real-time traffic without any explicit changes to the Internet service model. The essence of real-time service is the requirement for some service guarantees, and we argue that guarantees cannot be achieved without reservations. … We conclude that there is an inescapable requirement for routers to be able to reserve resources, in order to provide special QoS for specific user packet streams, or 'flows'.").

13 *See, e.g.,* Richard T.B. Ma, et al, *On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers*, Procs of 2008 ACM Conf Emerging Network Experiment and Tech (CoNEXT 2008), Madrid, Spain, December, 2008 ("we find the justification of the existence of paid-peering between transit ISPs. ... Our previous work ... showed that ... selfish ISPs have incentives to perform globally optimal routing and interconnecting decisions to reach an equilibrium that maximizes both individual profit and global social welfare. ... In this paper we extend our model ... Our result [finds instances where paid-peering can benefit welfare].") David Clark, *Network Neutrality: Words of Power and 800-Pound Gorillas*, 1 Int'l J. Comm. 701 (2007) ("As a technical mechanism, QoS seems to be beneficial. It directly addresses the real performance requirements of different sorts of Internet traffic … This reality begs the question of whether we can find a set of rules that might distinguish between "good" or "acceptable" forms of discrimination, and "bad" discrimination. Unless we can find a bright line, using regulation of discrimination to define acceptable behavior may cause more trouble than it cures.").

conduct already proscribed by antitrust law, such that further regulation by the Commission would largely be duplicative and generally unwarranted.

Nonetheless, the economic literature of vertical integration provides useful guidance for the Commission to consider. As with the literature on open platforms and innovation, the conclusions relating to vertical integration are consistently ambiguous. Most vertical integration *can* harm consumers; but most vertical integration also *can* benefit consumers. Importantly, in practice firms generally vertically integrate to capture efficiencies, and in so doing such integration ultimately benefits consumers.[14] As explained by FTC Commissioner and Professor of Law and Economics Josh Wright,

> Over a century of antitrust jurisprudence, economic study, and enforcement agency practice have produced a well-understood economic analysis of the competitive effects of a vertically integrated firm's "discrimination" in favor of its own products or services, including widespread recognition that such arrangements generally produce significant benefits for consumers.[15]

One of the most illustrative examples of misplaced concerns about vertical integration leading to harmful regulatory intervention is the Supreme Court's 1948 *Paramount* case, in which the Court broke up the vertically-integrated movie production and distribution industry.[16] Subsequent decades of analysis have not been kind to this action, finding that the vertical disintegration led to, *inter alia*, higher ticket prices for consumers, reduced variety and quality of films being produced, substantial loss of jobs within the movie industry, and, generally, the end of the golden age of Hollywood and the beginning of a multi-decade dark age.[17] One of the central reasons for this is that vertical integration allowed firms to better absorb risk and to better capture the returns on risky investment; conversely, the vertical disintegrated firms faced greater exposure to risk and less ability to capture the rewards of risky investment. The result was a substantial reduction in innovation in the industry: vertical integration can be a key ingredient for basic research, development, and innovation.

II.c. Multi-sided markets

A final related area of literature relates to multi-sided markets: markets in which two or more groups of users interact with one another by means of some platform. The Internet is a prototypical example of a multisided market, with ISPs and backbone providers acting as platforms that facilitate the interaction between end users and edge providers. To not needlessly belabor the point, as with the prior examples, the relevant literature yields consistently inconsistent results. In multisided markets, prohibiting platforms from engaging in discriminatory pricing can yield either consumer benefits or harm.[18]

---

14 *See, e.g.,* James C. Cooper et al., *Vertical Antitrust Policy as a Problem of Inference*, 23 Int'l J. Indus. Org. 639, 658 (2005) (surveying the literature and finding the vast majority of studies find vertical integration to have procompetitive effects).

15 Joshua Wright, *Defining and Measuring Search Bias: Some Preliminary Evidence,* Geo. Mason L. & Econ. Res. Paper. No. 12-14, at 5 (2011).

16 *U.S. v. Paramount Pictures, Inc.*, 334 U.S. 131 (1948).

17 *See generally* F. Andrew Hanssen, *Vertical Integration during the Hollywood Studio Era*, 53 J. L. and Econ. 519 (2010).

18 The literature here is voluminous, often demonstrates benefits from non-neutrality, and consistently notes ambiguous results. For some examples (most of which cite to the broader literature) see: Nicholas Economides and Joacim Tåg, *Network neutrality on the Internet: A two-sided market analysis*, 24 Information Economics and Policy 91 (2012) ("We have showed that one can find such parameter ranges both in the monopoly model and in the duopoly model suggesting that network

The examples discussed above make abundantly clear the indeterminacy inherent in any Open Internet rules. It is difficult, if not impossible, to know *ex ante* whether any given business model or type of conduct by online intermediaries will be beneficial or harmful to consumers. Any rules that mandate or foreclose certain practices or conduct on an *ex ante* basis would therefore be, of necessity, arbitrary or capricious. The converse, however, is also true: it is possible that certain practices or conduct can harm consumers, such that the Commission must develop an approach to the regulation of the Internet that allows it to take action where action is, in fact, necessary.

This suggests that the Commission should prefer to adopt general standards over clear rules. The legal basis for such an approach is discussed in the next Part.

Before turning to that discussion, it is useful to urge here that, should the Commission adopt clear rules it must affirmatively embrace and discuss the relevant economic literature to explain why a rule of general applicability is in the public interest. Given the uncertain consumer welfare effects that run throughout the relevant literature, the Commission must affirmatively address these concerns in order to avoid judicial challenge for any rules that it adopts. Indeed, such challenge will likely be warranted, given the arbitrariness such rules would demonstrate. More problematic, such a legal challenge would almost certainly result in further years of consumer doubt and uncertainty relating to network neutrality, which would cause concomitant consumer harm.

III. Legal authority and implementation

A central question raised by the NPRM is whether the FCC should act based upon Section 706 or Title II. This question has been overwrought: by and large, either approach gives the Commission broad and sufficient authority to enforce its proposed rules. There may be some differences on the margin as to what either approach will allow; but by and large, both

---

neutrality regulation could be warranted even when some competition is present in the platform market. *However, the overall effect of implementing network neutrality regulations can still be both positive and negative depending on parameter values.*") (emphasis added); Paul Njoroge, et al, *Investment in Two-Sided Markets and the Net Neutrality Debate*, 12 Review of Network Economics (Feb 2014) ("This paper adds to the growing body of formal economic analysis that will help inform policy makers on the net neutrality debate and sheds light on the validity, or lack thereof, of the arguments proposed by the different advocacy groups involved. In particular, this article develops a game theoretic model based on a two-sided market framework … to investigate the effects of a net neutrality mandate on investment incentives of ISPs, and its concomitant effects on social welfare, consumer and CP surplus, and CP market participation. ... More specifically, the results regarding the comparison between the neutral and non-neutral regimes for our theoretical and numerical-simulation models are as follows. *In both models, the non-neutral regime leads to a higher overall social welfare*. This result is driven by the higher investment levels caused by the non-neutral regime, which in turn increase consumer surplus and CP gross surplus.") (emphasis added); Jay Pil Choi, Byung-Cheol Kim, *Net Neutrality and Investment Incentives*, 41 RAND Journal of Economics (2010) ("Considering all three channels through which net neutrality can have an influence upon short-run total welfare, we can conclude that static welfare implications of net neutrality regulations depend on the trade-off between transportation cost saving and inefficient production. If the margin difference is significantly large relative to the degree of product differentiation, the discriminatory network would be preferred from the viewpoint of social welfare."; "We find that the relationship between the net neutrality regulation and investment incentives is subtle. Even though we cannot draw general unambiguous conclusions, we identified key effects that are expected to play important roles in the assessment of net neutrality regulations.").

approaches offer substantial, but not complete, authority. Given the broad grant of power either of these approaches affords, the more important question is how the Commission's choice of legal basis will impact its implementation of the rules. Implementation considerations strongly urge the Commission both to proceed based upon Section 706, and to implement its rules on a case-by-case enforcement model.

III.a. Section 706 vs. Title II

As upheld by the DC Circuit in January, the Commission has broad authority to regulate the Internet under Section 706.[19] The breadth of this power is largely in the Commission's own discretion, based on principles of deference.[20] The broad scope of this deference is long enduring and has been recently affirmed by the Supreme Court.[21] The primary limit is that the Commission cannot interpret Section 706 to do that which is expressly prohibited by the Communications Act – that is, it cannot regulate information services as common carriers.[22]

Title II is similarly broad. Title II grants the Commission pervasive authority to regulate broad aspects of the telecommunications sector, including myriad statutory sections authorizing specific regulation of prices, interconnection, and business practices – many of which clearly would not or should not fully apply in the Internet context. It is, however, not without limits. In specific, Title II is limited to various "unjust and unreasonable" practices.[23] As made clear previously, the various forms of conduct underlying the concerns the NPRM seeks to address have consistently ambiguous results, such that it is typically unclear, *ex ante*, whether any given conduct is in fact "unjust or unreasonable."

Importantly, clearly expressed statutory language – such as "unjust or unreasonable" – must be given meaning. Congress has very clearly expressed its intent that Title II does not apply to all conduct by telecommunications carriers that falls under Title II; rather, only the narrow class of conduct identifiable as "unjust or unreasonable" can be proscribed by the Commission. In the context of the NPRM, this is an exception that can quickly swallow the whole of Title II authority. Should the Commission choose to go down the Title II path, it will be little more than an invitation for further years of litigation. Curiously, should the Commission choose to adopt strong *ex ante* rules, this suggests that Section 706 may, in practice, give the Commission greater authority than Title II.

The Commission's forbearance authority is of little help in this context. The purpose of Section 10,[24] allowing forbearance of Title II's various requirements, is to allow the Commission to reduce the requirements of Title II in response to changing market conditions – it is not to allow the

---

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

20 *Chevron U.S. A. Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837 (1984).

21 *Arlington v. FCC*, 569 U.S. ___ (2013).

22 *Verizon* ("Given that the Commission has chosen to classify broadband providers in a manner that exempts them from treatment as common carriers, the Communications Act expressly prohibits the Commission from nonetheless regulating them as such.")

23 47 USC 202(a) reads in full:

> It shall be unlawful for any common carrier to make any *unjust or unreasonable discrimination* in charges, practices, classifications, regulations, facilities, or services for or in connection with like communication service, directly or indirectly, by any means or device, or to make or give any *undue or unreasonable preference or advantage* to any particular person, class of persons, or locality, or to subject any particular person, class of persons, or locality to any *undue or unreasonable prejudice or disadvantage*. (emphasis added)

24 47 USC 160.

Commission to craft an entirely new regulatory regime. Forbearance was adopted as part of the 1996 Telecom Act's goal of fostering competition in the telecommunications market, recognizing that the justifications for Title II's extensive regulation would diminish over time as entry occurred – and that in some cases Title II could unduly burden potential entrants.[25]

The Supreme Court's recent decision in *Utility Air Group v. EPA* should give the Commission great pause before proceeding down the Title II path.[26] In *Utility Air Group*, the EPA adopted broad regulations over greenhouse gases. In so doing, it acknowledged that the adopted rules were overbroad and administratively infeasible. In order to enforce the new rules as the agency intended, the EPA adopted a "Tailoring Rule" to tailor its enforcement of the regulations to only apply to the conduct the agency intended to target. The Supreme Court rejected this approach, explaining that "agencies must operate within the bounds of reasonable interpretation" and that "an agency interpretation that is inconsistent with the design and structure of the statute as a whole does not merit deference."[27] Like the EPA's regulation of greenhouse gasses, Title II regulation is meant to enable pervasive regulation of clearly problematic conduct, and the clear Congressional intent is that the forbearance mechanism is meant to foster conduct that will render the purposes of Title II nugatory. The Commission could not give meaning to Title II's "unjust and unreasonable" limitations without doing irreparable damage to the underlying statutory structure. The use of forbearance as a "tailoring rule" to structure Title II to suit the Commission's newfound purpose would be "inconsistent with the design and structure of the statute as a whole."

III.b. Implementation of legal authority

No matter whether the Commission acts under Section 706 or Title II, any strong *ex ante* rules that it adopts are likely to be challenged. This follows from the inherent uncertainty as to the effects of any proscribed conduct. Such challenge is particularly likely should the Commission act under Title II, given the concerns described above.

In deciding how to implement its rules the Commission must understand that both the industry and public need resolution over the Commission's Open Internet rules. Ongoing uncertainty over these rules constitutes a real harm to both consumers and industry. And, while telecommunications lawyers enjoy the likely hundreds of millions of dollars that have been spent on net neutrality over the past decade, few likely think this has been money well spent.

The basic problem with adopting strong *ex ante* rules is that whether any specific business model or conduct is harmful is highly fact-specific. Any rules addressing such conduct generally and on prophylactic terms are, therefore, by definition arbitrary and capricious. Importantly, administrative law has developed longstanding principles to address precisely this concern – chief among them agency choice of procedure. It is a longstanding and bedrock principle of administrative law that agencies that have the authority to develop rules of general applicability can develop those rules through case-by-case adjudication. As explained in the Supreme Court's seminal *Chenery II* opinion:

---

25 *See, e.g.,* Senate Report No. 103-367 (1994) ("S.1822 gives the FCC greater regulatory flexibility by permitting the FCC to forbear from regulating carriers when it is in the public interest. This provision will allow the FCC to reduce the regulatory burdens on new entrants. It will also permit the FCC to reduce the regulatory burdens on the telephone company when competition develops or when the FCC determines that relaxed regulation is in the public interest.")

26 *Util. Air Regulatory Grp. v. EPA*, ___ U.S. ___ (U.S. June 23, 2014)

27 *Id.*, slip op. at 10.

> Not every principle essential to the effective administration of a statute can or should be cast immediately into the mold of a general rule. Some principles must await their own development, while others must be adjusted to meet particular, unforeseeable situations. In performing its important functions in these respects, therefore, an administrative agency must be equipped to act either by general rule or by individual order. To insist upon one form of action to the exclusion of the other is to exalt form over necessity.
>
> In other words, problems may arise in a case which the administrative agency could not reasonably foresee, problems which must be solved despite the absence of a relevant general rule. Or the agency may not have had sufficient experience with a particular problem to warrant rigidifying its tentative judgment into a hard and fast rule. Or the problem may be so specialized and varying in nature as to be impossible of capture within the boundaries of a general rule. In those situations, the agency must retain power to deal with the problems on a case-to-case basis if the administrative process is to be effective. There is thus a very definite place for the case-by-case evolution of statutory standards.[28]

The Commission's proposed "commercial reasonableness" standard follows precisely this model – and is both learned and wise. Among the reasons for this are that the Commission clearly does have broad enforcement authority, such that a standards-based approach is less subject to judicial challenge outside of the context of a challenge to actual potentially-harmful conduct, and that any challenge to such conduct will be focused on whether that conduct actually is harmful. If such challenged conduct is actually harmful, the fact of that harm will go a long way toward demonstrating the legitimacy of the Commission's rules. If such challenged conduct is not actually harmful, a judicial loss will largely be cabined, affecting only the specific conduct in question and not jeopardizing the Commission's overall framework.

Should the Commission choose to proceed on a case-by-case basis, it should nonetheless issue general guidelines that provide notice as to the sort of conduct that may or may not raise concerns and the legal basis and metrics by which those concerns will be evaluated. The current proposed commercial reasonableness standard does an admirable job on this front.

Providing such guidelines is necessary to meet Constitutional Due Process and Fair Notice requirements. As explained in the Supreme Court's still recent *Fox II* opinion, "A fundamental principle in our legal system is that laws which regulate persons or entities must give fair notice of conduct that is forbidden or required."[29] Importantly, Constitutional problems arise not "because it may at times be difficult to prove an incriminating fact but rather because it is unclear as to what fact must be proved."[30] In other words, the difficulty of demonstrating that a given business model or conduct is problematic does not mean that the Commission cannot proscribe such conduct, but if the Commission is to proscribe that conduct it must be clear as to how it will evaluate whether the conduct is problematic.

---

[28] *SEC v. Chenery Corp.*, 332 U.S. 194, 202 (1947).
[29] *FCC v. Fox Television Stations, Inc.*, 132 S. Ct. 2307 (2012)
[30] *Id*.

IV. Statistical multiplexing, discrimination, and minimum levels of access

The final subject that these comments address is the relationship between the proposed no-blocking and nondiscrimination rules, and in particular, how prioritization affects "best effort" services and the minimum level of access required for a service not to be blocked. These issues are addressed from a technical perspective, focusing on statistical multiplexing – the basic means by which multiple users and applications are able to share Internet facilities. Importantly, many advocates for strong Open Internet rules assert that prioritization necessarily means a reduction in capacity available for non-prioritized services. From a technical perspective this understanding is simply not correct. This section explains why and offers other policy implications derived from the technical mechanisms controlling how data is sent over the Internet.

IV.a. A brief primer: statistical multiplexing, congestion, and queue management[31]

The Internet is a shared facility. The basic mechanism by which the Internet works is that human-understandable information that is to be communicated between endpoints is digitized and broken into small packets. These packets are sent over a shared communications infrastructure that routes them from a sender to a receiver; the receiver reassembles the packets and converts them back into human-understandable form. This is generally called a "packet switching" network, referring to the use of packetized data to communicate between endpoints.

But the packetization of data isn't the only basic feature of the Internet. Once data is packetized, it needs to be communicated across the network. This is done by a process known as statistical multiplexing. The basic idea behind statistical multiplexing is that a significant portion of any communications channel goes unused by any given application. For instance, a typical voice conversation consists of more silence (gaps between sounds made by the speakers) than speech. Other applications can make opportunistic use of that capacity without adversely affecting the initial application. Statistical multiplexing is the mechanism by which unused capacity by one communications channel is shared by other communications channels. This has two general benefits. First, without statistical multiplexing, five "conversations" would require five communications channels; with statistical multiplexing, those five conversations can be accommodated by two, or perhaps three, communications channels. This reduces cost and complexity of a communications network. And second, this allows individual "conversations" access to much greater capacity during periods of relative quiet on the network.

Packetization and statistical multiplexing are closely related: packet switching is the mechanism that makes statistical multiplexing possible on the Internet. Prior to the advent of packet switching, telecommunications services used either circuit-switched or dedicated lines. Under this model, lines were allocated to specific users for specific periods of time. As a result, they could not be shared between other simultaneous users.[32] Packet switching changed this, allowing for a much finer-grained level of sharing, on a purely statistical basis.

The benefits of statistical multiplexing are immense. Because most communication is "bursty," few applications need the full capacity of a dedicated communications channel.

---

31 The discussion that follows is meant to give a general sense of complicated technical concepts, to demonstrate their importance to the Open Internet discussion. Some liberties have been taken as to precise technical descriptions in order to make the discussion reasonably accessible to a generalist audience.

32 This description provides a gloss on the concept of statistical multiplexing. Technically, circuit-switching also is a form of statistical multiplexing, which connections are multiplexed on a much coarser boundary.

Statistical multiplexing therefore allows substantially more aggregate capacity to be offered to a group of users than would otherwise be possible if each user was allocated a dedicated communications channel. Conversely, any additional capacity added to a network to benefit one or a small group of users also creates additional capacity for all other users. Thus, because statistical multiplexing allows users to opportunistically use other users' unused capacity, incremental capacity benefits all users.

Statistical multiplexing is implemented primarily by two mechanisms: the algorithms that individual computers use to determine the rate at which they inject packets into the network ("congestion control" algorithms), and the algorithms that routers use to sort and prioritize packets ("queuing disciplines").[33] Queuing disciplines are particularly important. As routers receive packets from multiple sources they store those packets in temporary buffers before processing them to be sent to their destination. Early routers would process packets in the order that they arrived (i.e., on a first-in-first-out ("FIFO") basis) (fig. 1); if buffers were full, the router would drop any incoming packets (i.e., on a "tail-drop" basis). On naïve inspection, this appears to be a fair (even "neutral"!) approach to handling traffic. In practice, however, FIFO and drop-tail algorithms can result in some streams getting a disproportionate share of available capacity, other streams being locked out of getting any capacity at all, and overall inefficient levels of network utilization.

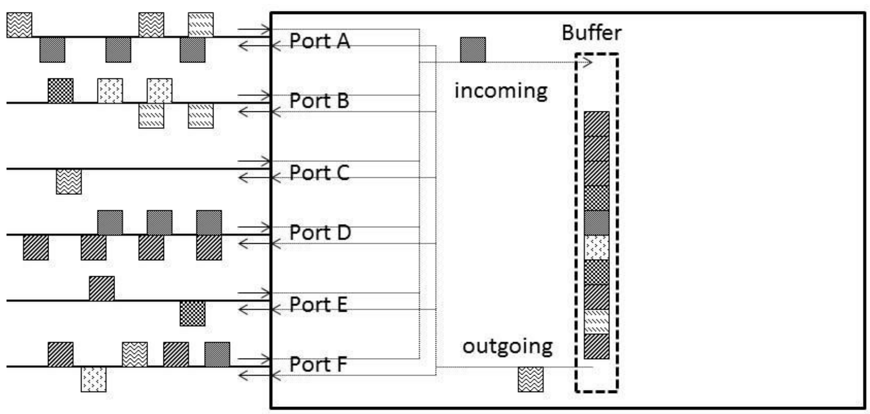

A basic router. As packets arrive they are added to a buffer for processing. The router processes packets one at a time, sending them to a destination port, in the order in which they arrived.

Fig. 1

Given these concerns, engineers have been developing various active queuing disciplines – generally referred to as Active Queue Management ("AQM") – for decades.[34] AQM algorithms allow routers to adopt different approaches to managing queues and dropping packets as buffers reach capacity, in order to ensure "fair"[35] and efficient network operation. This is the mechanism by which "paid prioritization" is most likely to be implemented: ISPs

33 *See generally* Braden, et. al., *RFC 2309: Internet Performance Recommendations* (April 1998), *available at* http://www.ietf.org/rfc/rfc2309.txt. The term "queuing discipline" used here encompasses both queue management and scheduling algorithms.

34 *See, e.g.,* Floyd, S., and Jacobson, V., *Random Early Detection gateways for Congestion Avoidance*, 1 IEEE/ACM Transactions on Networking 397 (1993).

35 "Fair" is a term of art, referring to how congestion control algorithms and router queuing disciplines interact to apportion bandwidth on a shared resource. There are multiple definitions, and no consensus definition. Generally, "fairness" means that, on a connection shared by *N* connections, no connection will prevent any other connection from being able to use at least 1/N of the overall capacity.

can configure their routers to handle prioritized data specially. For instance, prioritized traffic can be moved to the head of the queue. Or, more often, prioritized traffic will be placed in a separate queue, from which packets will be processed more often than from other queues (fig. 2.).[36]

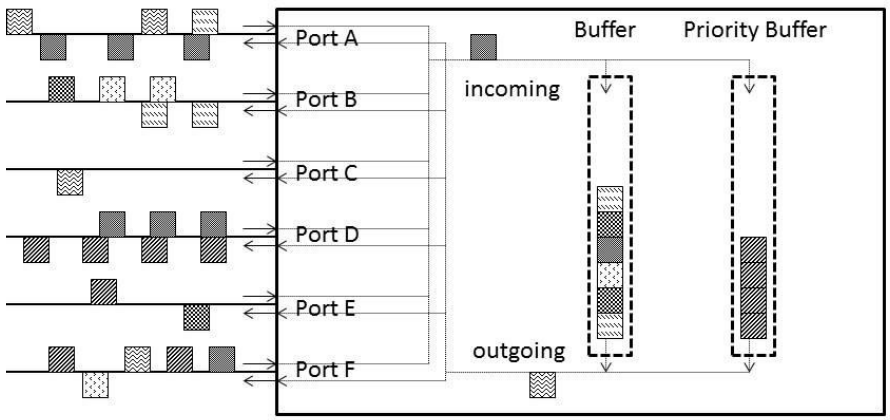

One model of a router with prioritization. As packets arrive, most go into one buffer, while "priority" packets go into their own buffer. The router can process packets from the two buffers in many ways, for instance by alternating between the buffers. This reduces the time that prioritized packets spend "waiting."

Fig. 2

Under any queuing discipline, congestion occurs in one of two ways: packets are either dropped, or they are delayed. Drops occur either when a router's buffer space is full, or when the network connection between devices is at capacity (e.g., trying to transmit data at a rate of 15 mbps for a sustained period over a connection only capable of transmitting 10 mbps).[37]

Delays occur when packets arrive at a router faster than the router can send them to their destination. Where this happens, the router buffers packets until it can process them. Such delays most often occur because Internet traffic is "bursty," meaning that packets from a given sender tend to arrive at a router in groups – not because the destination network has insufficient capacity to handle the data to be sent over it. As an example, if we assume a router with ten 100 mbps ports (and symmetric traffic patterns across all ports), with average

---

36 Importantly, queuing disciplines alone cannot efficiently guarantee an equal allocation of bandwidth among streams. Bandwidth allocation is governed by many factors, including congestion on other parts of the network, latency and packet round-trip-time, and congestion algorithm. Certain queuing disciples attempt to ensure equal bandwidth between streams; while such disciplines can increase the share of bandwidth allocated to lower-speed streams (e.g., those that take longer to reach a steady-state packet rate due to relatively higher round-trip times), they often accomplish this at the cost of disproportionately slowing other streams.

37 In practice, these cases are generally equivalent today. This is because most devices are on one of two types of connection. Either they are on switched ports, meaning that there is only a single sender/receiver pair on each connection. Or they are on an asymmetric shared connection with a single gateway (e.g., the CMTS). In that case, devices on the shared connection may interfere with one another, resulting in packets being sent to the gateway to drop. However, the bulk of data is typically being sent to the devices on the shared connection (e.g., from edge providers, through the CMTS, to end-users). Where this is the case, the gateway is the only transmitted, so drops are likely to result as data from edge providers arrives as a rate exceeding the capacity of the shared connection; in which case, the gateway will buffer packets until its buffers are exhausted, at which point it will begin dropping packets. The great exception to this discussion is wireless networks, where packet drops are more common, for various reasons.

utilization of 40 mbps per port, it is likely that the rate of incoming traffic to each port will regularly exceed several hundred mbps for very short periods of time. Without buffering, this would result in substantial packet loss (i.e., dropped packets); with buffering, the router can queue several hundred or thousand packets (or more) received over, e.g., milliseconds to 10 of milliseconds and send them back out over a period of, e.g., 10s of milliseconds to 100s of milliseconds. Most applications are not particularly sensitive to delays – and even those that are can generally tolerate ther moderate variances that are typically seen on actual networks – so it is often preferable to address congestion with buffering instead of by dropping packets.[38]

Another example is particularly illustrative, as it demonstrates a common instance where buffers may be overrun, resulting in packet loss. The MPEG-DASH protocol, which is used by services like Netflix to stream video with adaptive bitrates, sends video in chunks (commonly in units of 10 seconds of video each).[39] This typically results in an ON/OFF transmit pattern, where video is sent from the server to the client for about 10 seconds, and then sends nothing for about 10 seconds. Researchers are only just starting to understand how this protocol, which transmits for relatively large bursts, interacts with router buffering and TCP's congestion control algorithm. While on average data is being sent to a steady number of clients at a constant data rate, statistically there will be periods where either far more or far

---

[38] But note, this is not always the case – proper tuning of buffer sizes is a difficult and technical topic. *See, e.g.*, Jim Gettys & Kathleen Nichols, *Bufferbloat: Dark Buffers in the Internet*, 55 Comm's ACM 57 (2012). Additionally, while it may be preferable to address ongoing congestion by buffering instead of dropping packets, congestion *avoidance* is often best implemented by responding to delay. There is longstanding research and debate within the technical community over congestion control algorithms that rely on packet loss (e.g., TCP Reno) vs. packet delay (TCP Vegas). *See, e.g.,* Budzisz, et al., *On the Fair Coexistence of Loss- and Delay-Based TCP*, 19 IEEE/ACM Trans. on Networking 1811 (2011).

[39] For discussions of MPEG-DASH relevant to this and the next paragraph, see: Ahmed Mansy, et al., *SABRE: A client based technique for mitigating the buffer bloat effect of adaptive video flows*, Procs. 4th ACM Multimedia Systems Conf. 214 (2013) ("we show ... that a single DASH stream can cause significant delays to other ongoing applications sharing the home network in a typical residential setti"g."); Jim martin, et at., *Characterizing Netflix bandwidth consumption*, Procs. Consumer Comm's and Networking Conf 230 (2013) ("Ongoing academic research is providing foundations for understanding how DASH applications behave and how they might be improved."; "we seek to understand the impact of adaptive applications on congestion and bandwidth control mechanisms throughout the Internet or within a broadband access network. The work presented in this paper provides foundations for achieving this goal."; "Academic research on DASH is just emerging."; "The dynamics and implications of multiple levels of end-to-end congestion control are not well understood."); Saamer Akhshabi, et al., *An experimental evaluation of rate-adaptation algorithms in adaptive streaming over HTTP*, Procs. 2d ACM Multimedia Systems Conf. 214 (2013) ("Adaptive streaming over HTTP is a new technology. It is not yet clear whether the existing commercial players perform well, especially under dynamic network conditions. Further, the complex interactions between TCP's congestion control and the application's rate-adaptation mechanisms create a 'nested double feedback loop' - the dynamics of such interacting control systems can be notoriously complex and hard to predict."); Te-Yuan Huang, *Confused, timid, and unstable: picking a video streaming rate is hard*, Procs. 2012 ACM Conf on Internet measurement 225 (2012) ("all three [MPEG-DASH] services we study display degraded performance in the presence of competing traffic, well below the video quality possible if the client used its fair share of bandwidth."; "all three services we study display degraded performance in the presence of competing traffic, well below the video quality possible if the client used its fair share of bandwidth."; "In the worst case, the feedback loop creates a "death spiral" and brings the playback rate all the way down to its lowest value."); Junchen Jiang, et al., *Improving Fairness, Efficiency, and Stability in HTTP-based Adaptive Video Streaming with FESTIVE*, Procs. 8th Int'l Conf. Emerging Networking Experiments and Techs. 97 (2012) ("Many commercial video players rely on bitrate adaptation logic to adapt the bitrate in response to changing network conditions. Past measurement studies have identified issues with today's commercial players with respect to three key metrics---efficiency, fairness, and stability---when multiple bitrate-adaptive players share a bottleneck link. Unfortunately, our current understanding of why these effects occur and how they can be mitigated is quite limited.").

fewer than the average number of segments are being simultaneously sent. When more than the average number of segments are being simultaneously sent, given the size of each segment, there is a potential to saturate router buffers and cause packet-loss congestion. Such congestion may give the appearance that the network is under-provisioned by the network operator, when in fact it is largely the result of application-level attempts at congestion control; additionally, while it will adversely affect all traffic being sent over the shared link, it is likely to affect some traffic more adversely than other traffic, potentially giving the appearance of discriminatory treatment.[40]

Critically, from the perspective of the Commission's Open Internet efforts, this demonstrates the limits of our understanding of the basic technologies that the FCC is attempting to regulate. It is difficult to imagine that the Commission can implement an *ex ante* regulatory regime that yields an efficient allocation of these resources given the current state of knowledge about how these resources work. The interaction between MPEG-DASH, router queuing, and congestion control is a current area of cutting edge research. Current research suggests that traditional, best-effort, non-prioritized routing may yield substantially inefficient use of the network resource. It may well turn out to be the case that efficient routing of data like streaming video requires router-based prioritization. It may even turn out that efficient routing of streaming video data is necessarily harmful to other data – it may not be possible to implement a single network architecture that efficiently handles data with differentiated characteristics. If this is the case, then it may certainly be "commercially reasonable" that streaming video providers pay a premium for the efficient handling of their data, in order to compensate for the negative externalities that those uses impose upon other users and uses.

IV.b. Prioritization, Congestion, and Minimum Levels of Access

To date, no firm is known to have offered a paid-prioritization service of the sort contemplated by the NPRM. It is therefore not entirely clear what such a service would look like or how it would be implemented. While necessarily speculative – again demonstrating the danger of adopting strong *ex ante* rules – in all likelihood a paid-prioritization model would be implemented through priority buffering as described above: a firm would pay to have its packets handled on a priority basis by another firm's routers. Taking this as the basic mechanism of prioritization yields several important conclusions about the concepts of minimum levels of access and discrimination – and, importantly, their relationship to one another.

Prioritization can affect connections in one of two ways: by increasing or decreasing that connection's bandwidth, and by increasing or decreasing its delay. The primary effect of prioritization is to decrease a prioritized connection's delay; this follows from the basic prioritization mechanism of handling incoming packets via a priority queue. Unsurprisingly, this can increase the delay to which non-prioritized packets are subjected.

This increased delay[41] is unlikely to meaningfully affect most applications. While some applications are particularly sensitive to delay, delay is a persistent technical reality. All applications are subject to delay, and most are not sensitive to it; all delay-sensitive

---

[40] See, e.g., Te-Yuan Huang, *Confused, timid, and unstable: picking a video streaming rate is hard*, supra (describing the performance "death spiral" that service using MPEG-DASH (such as Netflix) can experience under congestion circumstances).

[41] This discussion does not differentiate between "delay" and "changes in delay" (technically known as jitter). For some applications, especially streaming media, jitter is more problematic than delay. Prioritization would likely yield similar effects for both delay and jitter, though the precise effects are ambiguous.

applications are designed to accept and mitigate delay as best they can. Any increase in delay caused by prioritization is likely to be marginal and fall within the accepted margins of typical delay-sensitive applications. Moreover, because delay is a persistent technical problem that is only exacerbated, not created, by prioritization, a better approach to delay mitigation is the wide-scale implementation of application-controlled or automated prioritization (e.g., QoS) mechanisms by routers. These mechanisms – many of which already exist, but are infrequently implemented – allow routers to prioritize delay-sensitive traffic. Wide-scale implementation of these mechanisms would broadly address any concerns raised by the adverse effects of paid-prioritization on delay-sensitive applications.[42]

Prioritization would also affect how much bandwidth competing streams are able to use. Assuming the network is uncongested – as recent research makes clear is the case on most consumer-Internet backbones – prioritization would allow a prioritized connection to consume more bandwidth. But, counterintuitively and contrary to the understanding of many net neutrality advocates, prioritization would not appreciably affect the bandwidth available to other connections.

This follows from the interaction of queuing disciplines and congestion control in statistically multiplexed systems. Almost every implementation of TCP in use today uses packet loss, not delay, to control its data rate.[43] End-hosts on the Internet do not know how much bandwidth is available to them, so they have no way of knowing at any given time what speed they should send data at.[44] What they do instead is start sending data at a slow speed and incrementally increase that speed until congestion occurs (e.g., packets are dropped). At that point, they know they have exceeded available capacity, so they slow down (usually by reducing transmit rate by half); they then resume incrementally increasing their speed. This process, known as "congestion control" repeats indefinitely.

Because congestion control algorithms primarily rely on packet-loss to signal congestion, so long as a network is not experiencing packet loss prioritization will not substantially affect the rate at which end-hosts send packets.[45] Thus, and directly contrary to

---

42 In considering the meaning of "minimum levels of access," the Commission may consider, in particular, whether Internet Service Providers implementing paid prioritization should also be required to implement user-selectable or automated QoS, perhaps with a requirement that packets subject to paid-prioritization cannot be given greater priority than applications reasonably recognized as especially delay sensitive. The merits of such a proposal require substantial further investigation, and these comments do not endorse such an approach without further research.

43 Standard implementations of TCP used by all major operating systems rely on the packet-loss mechanism. Delay-oriented congestion control algorithms also exist, but are less frequently used – largely because they are systematically unable to make as efficient use of the network resource in a system shared with packet-loss oriented TCPs. *See, e.g.,* Budzisz, et al., *On the Fair Coexistence of Loss- and Delay-Based TCP*, 19 IEEE/ACM Trans. on Networking 1811 (2011). The great exception to this is recent versions of Windows, which uses a hybrid loss- & delay-based algorithm ("Compound TCP"). Even under congestion, however, Compound TCP is bounded on the low end by the performance of the standard Reno loss-based algorithm.

44 This is a fundamental constraint of the Internet. Since no one can know, *ex ante*, how many users will be sharing a connection at a given time, one can never know how much bandwidth is available to a given user at a given time – even if all of the technical characteristics of the network are known (which they are not). For instance, if three users share a 100 mbps connection but only two are active at a given time, a "fair" TCP will let each of them use 50 mbps. If the third user then starts using the connection, only 33 mbps will be allocated to each. There is no way to know whether or when other users will be making use of the shared resource.

45 The rate at which TCPs increase the rate at which they send data is affected by delay, so there is some marginal effect on speed that results from prioritization. This is, however, a second order effect, so is bounded by a relatively small margin, even in the case of a linear increment. Modern TCPs increasingly use

claims of many advocates, prioritization of some connections does not reduce the capacity available to other connections. Or, if we use the rhetoric of the debate, on connections without substantial packet loss the hypothetical addition of "fast lanes" (prioritized connections) does not relegate other connections to a "slow lane."

IV.c. Implications

This discussion suggests several conclusions relating to the NPRM's consideration of minimum levels of access and prioritization, and in particular how the two relate.

First and foremost, the effects of prioritization should be measured primarily in terms of on-net packet loss, and, in particular, the relative levels of packet loss between prioritized and non-prioritized traffic. If these rates are similar (including the case where they are at or near zero), this strongly suggests that any prioritization is commercially reasonable. Indeed, it strongly suggests that prioritization is not affecting non-prioritized traffic at all. In cases where there is substantial packet loss but the rates are comparable between prioritized and non-prioritized traffic, this suggests the relevant network is under-provisioned – but it does not suggest concerns relating to prioritization or Open Internet principles generally.

The only situation where prioritization on uncongested networks may reasonably be seen as adversely affecting non-prioritized connections is where those connections are substantially delay-sensitive. In such cases, network providers implementing paid prioritization should take steps to ensure such applications are not unduly affected. Such steps could take many forms, from user-selectable prioritization, to automated (e.g., DPI-based) prioritization of delay-sensitive traffic, to other router-based QoS mechanisms (e.g., tuning the rate at which prioritized buffers are serviced relative to non-prioritized buffers). These and similar ideas have long been discussed within the technical community as important to efficient network use; paid prioritization may reasonably increase the urgency of their deployment. Given the (small, but hard to define) range of possible applications that could be adversely affected by prioritization, the uncertain ways and extent to which they may actually be affected in any particular case, and the myriad approaches to mitigating such effects, the Commission should address any concerns arising from prioritization on a case-by-case basis.

In all other cases, the fact that a network is uncongested strongly suggests both that any prioritization is commercially reasonable, and that such prioritization is not adversely affecting other connections (e.g., any minimum level of access that would be provided without prioritization is being maintained on the network with prioritization).

Because a great portion of a given connection's performance is determined by circumstances outside the control of a given network operator (i.e., by off-net factors), competitive benchmarking of services' on-net performance is an important tool in understanding how, or whether, any on-net prioritization is affecting a given service's performance. For instance, as discussed above, the performance characteristics of MPEG-DASH, the protocol currently used by firms like Netflix to deliver streaming video, are not yet well understood. Initial research suggests that MPEG-DASH in particular has poor performance characteristics that are attributable to its design and interaction with TCP's congestion control mechanisms – and that are outside of the control of a given network operator.

---

non-linear recovery mechanisms (e.g., CUBIC, New Reno, Compound TCP), for which a marginal increase in delay is unlikely to substantially affect average transmission rates.

Operationalizing this factor, in evaluating whether prioritization unreasonably affects a given service the Commission should look to the performance of other similar services. If other similar services are able to operate well on a given network, this strongly suggests that any performance issues are related to off-net factors. Even in the absence of such positive comparisons, the Commission should be careful to look at the underlying technical mechanisms and state of relevant research in attributing any performance difficulties to on-net vs. off-net factors.

Related to this point, the earlier discussion of statistical multiplexing leads to a final observation: services that consume a statistically disproportionate amount of the shared network resource (e.g., available capacity) impose a negative externality upon all other users and uses of that resource. Independent from any fees for prioritization, such users should reasonably be expected to pay a premium for their disproportionate use of the resource. A basic premise of the Internet is that aggregation of underutilized resources creates a more efficient aggregate resource. This aggregation is accomplished by statistical multiplexing. But when a single service uses a disproportionate share of the network resource, the opportunistic sharing enabled by statistical multiplexing (i.e., allowing one service to use another service's excess capacity) becomes free-riding instead. Such conduct is more harmful to the Internet's future as an innovative platform than prioritization or other concerns raised in the NPRM; to whatever extent the Commission takes action to police conduct such as prioritization, similar attention should be paid to services that make uncompensated disproportionate use of the network resource.

Conclusion

These comments are respectfully submitted, with the purpose of making clear the need for the Commission to proceed on this matter (to the extent that it proceeds at all) on a case-by-case basis, using administrative adjudication and enforcement actions to address specific conduct in a fact-specific setting. Such an approach is necessary because at an economic and technical level the concerns that have given rise to the Open Internet proceedings have the potential to either harm or benefit consumers. The effects of any particular form of conduct generally cannot be evaluated on an *ex ante* basis, but rather must be evaluated in the context of an actual case or controversy. As such, any strong *ex ante* rule would necessarily be arbitrary and capricious, and otherwise in contravention of modern principles of administrative law and in excess of the Commission's legal authority. Rather, case-by-case adjudication is the typical and preferred approach for Administrative agencies to use in developing legal norms in settings such as this proceeding. Finally, any action that the Commission does undertake, either on a rulemaking or adjudicatory basis, must be based on a sound understanding of the underlying technical principles on which the Internet operates, and with substantial deference to the current state of technical knowledge. Much of the discussion relating to this proceeding, including the basis for many of the underlying concerns, either misunderstands or disregards the underlying technology, often drawing inapt (or inept) conclusions or raising technically implausible concerns. This includes concerns raised by many individuals and firms who work in technology-related fields, but who lack significant expertise relevant to these proceedings. One of the great achievements of the Internet is that one need not understand the vagaries of congestion control or queuing disciplines to create a cutting edge technology firm. Any action undertaken by the Commission must be based, first and foremost, on a sound understanding of the underlying technology, lest the Commission risk both acting in contravention of its legal mandate and significantly undermining the efficient operation of the Internet.