On the Relationship Between QoS & QoE:

Why Differential Traffic Management on the Internet Is Not a Zero-Sum Practice

*****

A Research Paper for TPRC 44: The Research Conference On Communications, Information and Internet Policy

by Thomas W. Struble[1]

August 31st, 2016

---

[1] Thomas W. Struble is Policy Counsel at **TechFreedom**, a non-profit, non-partisan technology policy think tank based in Washington, DC. Struble obtained his Juris Doctorate from the George Washington University Law School, where he was a Presidential Merit Scholar and an editor on the Federal Communications Law Journal. Struble also holds Bachelor degrees in Psychology and Political Science from the University of Kansas, from which he graduated with the highest distinction.

This work builds upon the work of many others who have come before, with the ultimate goal of advancing scholarship and understanding in the developing area of Internet traffic management and telecommunications policy. A great deal of gratitude is owed to Berin Szóka, Ben Sperry, Alex Kreilein, and Ashkhen Kazaryan for their collaboration and assistance in helping to transform a rather notional topic into this final paper, and leading the author on a great voyage of discovery along the way.

Abstract

Is differential traffic management (i.e., prioritizing some forms of Internet traffic over others) a zero-sum practice? Do the benefits of favoring Internet traffic from some edge providers come only at the expense of other edge providers? Or, rather, do the benefits of prioritization exceed the corresponding harms, making this a positive-sum practice that enhances aggregate welfare?

In paragraphs 125–32 of the 2015 Open Internet Order, the FCC laid out its case for why "paid prioritization" — defined as "the management of a broadband provider's network to directly or indirectly favor some traffic over other traffic, including through use of techniques such as traffic shaping, prioritization, resource reservation, or other forms of preferential traffic management, either (a) in exchange for consideration (monetary or otherwise) from a third party, or (b) to benefit an affiliated entity" — should be banned, ex ante, as a per se unreasonable practice.

To justify this proscription, the FCC pointed to support in the record from commenters asserting that differential traffic management is "inherently a zero-sum practice," which will inevitably lead to a bifurcated Internet wherein certain edge providers are relegated to a "slow lane" with degraded performance. The FCC also cited to multiple opposing comments in the record that argued traffic prioritization is not a zero-sum game, but, evidently, did not find these arguments persuasive.

Drawing on principles of game theory, cognitive psychology, and statistics, this paper examines this discrete issue and posits that — contrary to the arguments accepted by the Federal Communications Commission — differential traffic management on the Internet is not a zero-sum practice. Specifically, this paper recognizes the relationship between the quality of service ("QoS") for edge services and quality of experience ("QoE") for users, and observes that while these two variables are positively correlated for all edge services, the strength of such correlation differs significantly. For edge services that are particularly sensitive to latency, packet loss, or bandwidth constraints (e.g., live video, online multiplayer gaming), the users' QoE will have a strong direct correlation to the services' QoS. Conversely, for edge services that are comparatively insensitive to latency, packet loss, and bandwidth constraints (e.g., email, software updates), the direct correlation between QoS and QoE will be significantly weaker.

Thus, if an ISP were to prioritize services that fall into the former category, and de-prioritize services that fall into the latter category, the benefits yielded to the former (i.e., higher user QoE) would outweigh the harms done to the latter (i.e., slightly lower user QoE), so the network as a whole would work better and aggregate consumer welfare would be increased. For example, in such a system, users streaming live video would be much happier, while users checking their email would be only a little worse off, and may see no significant reduction in QoE. This suggests that the FCC should perhaps reconsider its view of such traffic management practices going forward, particularly as new services with greater QoS demands (e.g., VR-based services) begin to enter the market.

Table of Contents

Introduction

> Governments of the Industrial World, you weary giants of flesh and steel . . . .
>
> . . . .
>
> You claim that there are problems among us that you need to solve. You use this claim to invade our precincts. Many of these problems don't exist. . . .
>
> . . . .
>
> In China, Germany, France, Russia, Singapore, Italy, and the United States, you are trying to ward off the virus of liberty by erecting guard posts at the frontiers of Cyberspace. These may keep out the contagion for a small time, but they will not work in a world that will soon be blanketed in bit-bearing media.[2]

The Internet is a triumph of human spirit and ingenuity. It "was created by — and continues to be shaped by — decentralized groups of scientists and programmers and hobbyists (and more than a few entrepreneurs) freely sharing the fruits of their intellectual labor with the entire world."[3] Truly, the State played a key role in funding some of the initial work that built the foundation for the modern Internet,[4] and the Internet still reflects this State influence,[5] but, by and large, the success story of the Internet is one of bottom-up emergent progress, rather than top-down directed progress. "Nobody is in charge. Yet for all its messiness, the internet is not chaotic. It is ordered, complex, and patterned. It is a living example, before our eyes, of the phenomenon of evolutionary emergence — of complexity and order spontaneously created in a decentralised fashion without a designer."[6]

Nevertheless, as Leviathan is wont to do, the State has increasingly found it appropriate to intervene in the digital realm and regulate certain facets of the Internet. Most recently, the Federal Communications Commission ("FCC") imposed a broad set of new regulations on

---

[2] John Perry Barlow, *A Declaration of the Independence of Cyberspace*, ELEC. FRONTIER FOUND. (Feb. 8, 1996) (speech delivered in Davos, Switzerland), https://goo.gl/Y9gnvJ.

[3] Tom Gauld, *The Internet? We Built That*, NEW YORK TIMES (Sept. 21, 2012), *available at* http://goo.gl/PPrzv.

[4] *See, e.g.*, *id.*

[5] Indeed, the concept of packet switching and the Internet's distributed architecture were originally designed to enable it to withstand significant network outages — such as from a series of Soviet missile strikes — and still function. *See, e.g.*, *id.*

[6] MATT RIDLEY, THE EVOLUTION OF EVERYTHING: HOW NEW IDEAS EMERGE 412 (2015).

the Internet, in the 2015 Open Internet Order,[7] designed to police potential malfeasance by Internet service providers ("ISPs") and safeguard the principle of "net neutrality." Among the potential non-neutral ISP behaviors that concerned the FCC was something called "paid prioritization," which was defined as

> the management of a broadband provider's network to directly or indirectly favor some traffic over other traffic, including through use of techniques such as traffic shaping, prioritization, resource reservation, or other forms of preferential traffic management, either (a) in exchange for consideration (monetary or otherwise) from a third party, or (b) to benefit an affiliated entity.[8]

Interestingly, unlike the rules against blocking and throttling, the ban on paid prioritization does not contain an exception for reasonable network management, because, in the words of the FCC, "paid prioritization is inherently a business practice rather than a network management practice."[9] Given that ISPs are in the business of managing Internet traffic flowing over their networks, the contrast here between business practices and network management practices appears to be a distinction without a difference.

Stated plainly, the FCC's ban on paid prioritization limits the ability of an ISP to engage in differential traffic management (i.e., prioritizing some types of Internet traffic over others).[10]

---

[7] Protecting and Promoting the Open Internet, *Report and Order on Remand, Declaratory Ruling, and Order*, GN Docket No. 14-28, 30 FCC Rcd. 5601 (Feb. 26, 2015) ["2015 OIO"], *available at* https://goo.gl/dlx07R.

[8] 2015 OIO, ¶ 18.

[9] *See id*. ¶ 18 n.18.

[10] Throughout this paper, differential traffic management will be compared with undifferentiated, or "best-efforts," traffic management. The latter treats all Internet packets the same, whereas the former utilizes certain protocols that mark packets (in the IPv4/IPv6 packet header DiffServ field) and assign them a different set of transmission rules (called "per-hop behaviors") based on their particular needs, and treats them accordingly. For example, using best efforts traffic management, if 10 GBs of data are trying to pass through a 9 GB connection at any given point in time (e.g., one second), the ISP will drop every 10th packet until the congestion subsides. By contrast, an ISP using differential traffic management will give relative priority to traffic that is sensitive to latency or packet loss, and drop fewer of those packets or reduce latency for those services at the expense of lower priority traffic, which is less sensitive to latency and packet loss. *See generally* Nichols K. et al., *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers* (The Internet Society, Dec. 1998), *available at* https://goo.gl/P7mMMj (defining the use of the DS field in packet headers and a base set of packet forwarding treatments, or per-hop behaviors); *see also* Broadband Internet Tech. Advisory Grp., *Differentiated Treatment of Internet Traffic: A Uniform Agreement Report*, at i (Oct. 2015) ["2015 BITAG Report"], *available at* https://goo.gl/pS8aIb ("In its broadest sense, traffic differentiation includes any

Under the FCC's rules, differential traffic management is allowable, but only in certain forms,[11] and only under certain business arrangements,[12] which are subject to close scrutiny from the agency,[13] reflecting its view that these network practices allegedly "harm consumers, competition, and innovation, as well as create disincentives to promote broadband deployment[.]"[14] To justify this proscription, the FCC pointed to support in the record[15] from commenters asserting that differential traffic management is "inherently a zero-sum practice,"[16] which will inevitably lead to a bifurcated Internet wherein certain edge providers are relegated to a "slow lane" with degraded performance.[17]

Drawing on principles of game theory, cognitive psychology, and statistics, this paper examines this discrete issue and posits that — contrary to the arguments accepted by the

---

technique that classifies and applies potentially different treatment to two or more traffic flows contending for resources on a network (a flow being a group of packets that share a common set of properties).").

[11] *See* 2015 OIO, ¶ 124 n.281 ("While not within the definition of 'throttling' for purposes of our no-throttling rule, the slowing of subscribers' content on an application agnostic basis, including as an element of subscribers' purchased service plans, will be evaluated under the transparency rule and the no unreasonable disadvantage/interference standard."); *id* ¶ 125 n.284 ("Other forms of traffic prioritization, including those that serve a public safety purpose, may be acceptable under our rules as reasonable network management.").

[12] *See id.* ¶¶ 207–13 (describing the FCC's approach to "Non-BIAS Data Services," formerly referred to as "specialized services," such as IPTV and facilities-based VoIP, which may be offered by ISPs — in addition to broadband Internet access services — and are beyond the scope of the Open Internet Order's rules, but may still be deemed illegal if the FCC deems they are being used to circumvent the rules or are "undermining investment, innovation, competition, and end-user benefits").

[13] *See id.* ¶ 129 ("Given the potential harms to the virtuous cycle, we believe it is more appropriate to impose an *ex ante* ban on [paid prioritization], while entertaining waiver requests under exceptional circumstances."); *id.* ¶ 132 ("An applicant seeking waiver relief under this rule faces a high bar. We anticipate grating such relief only in exceptional cases.").

[14] *Id.* ¶ 125.

[15] *Id.* ¶ 19 (citing *Protecting and Promoting the Open Internet*, Comments of Mozilla, GN Docket No. 14-28, at 20 (July 15, 2014) ["Mozilla Comments"], *available at* https://goo.gl/ZRkdhA) ("The record demonstrates the need for strong action. . . . Mozilla, among many other such commenters, explained that '[p]rioritization . . . inherently creates fast and slow lanes.'").

[16] *See id*. ¶ 126 n.287 (citing Mozilla Comments, at 20).

[17] *Id.* ¶ 18 ("To protect against "fast lanes," this Order adopts a rule that [bans paid prioritization].").

Federal Communications Commission — differential traffic management on the Internet is not a zero-sum practice. Specifically, this paper recognizes the relationship between the quality of service ("QoS")[18] for edge services and quality of experience ("QoE")[19] for users, and observes that while these two variables are positively correlated for all edge services, the strength of such correlation differs significantly. For edge services that are particularly sensitive to latency, packet loss, or bandwidth constraints (e.g., live video, online multiplayer gaming), the users' QoE will have a strong direct correlation to the services' QoS. Conversely, for edge services that are comparatively insensitive to latency, packet loss, and bandwidth constraints (e.g., email, software updates), the direct correlation between QoS and QoE will be significantly weaker.

Thus, if an ISP were to prioritize services that fall into the former category, and de-prioritize services that fall into the latter category, the benefits yielded to the former (i.e., higher user QoE) would outweigh the harms done to the latter (i.e., slightly lower user QoE), so the network as a whole would work better and aggregate consumer welfare would be increased. For example, in such a system, users streaming live video would be much happier, while users checking their email would be only a little worse off, and may see no significant reduction in QoE. This suggests that the FCC should perhaps reconsider its view of such traffic management practices going forward, particularly as new services with greater QoS demands (e.g., VR-based services) begin to enter the market.

---

[18] Quality of Service, or QoS, is a shorthand term used to represent all the characteristics of a network connection, such as bit rate, latency, packet loss, jitter, and so forth. Different services require different levels of QoS in order to function. *See, e.g.*, Mario Marchese, QOS OVER HETEROGENEOUS NETWORKS (Wiley, 2007), *available at* http://goo.gl/vijatY ("The term 'QoS' is used in different meanings, ranging from the users' perceptions of the service to a set of connection parameters necessary to achieve particular service quality. The QoS meaning changes, depending on the application field and on the scientific scope."). The term is necessarily imprecise, but it is still useful for discussions such as the one in this paper.

[19] Quality of Experience, or QoE, is a shorthand term used to represent a user's perception of a particular service, typically in a holistic way that assesses the overall level of satisfaction the user had in using the service. *See, e.g.*, Int'l Telecomm. Union, ITU-T, Telecomm. Standardization Sector of ITU, Focus Group on IPTV, *Liaison Statement: Definition of Quality of Experience (QoE)*, at 2 (Feb. 6, 2007), *available at* https://goo.gl/CQ9TUy (defining QoE as "The overall acceptability of an application or service, as perceived subjectively by the end-user After briefly recounting the two successful legal challenges to previous FCC actions, the paper will then dissect the 2015 Open Internet Order and comment briefly upon its legal prospects before turning in earnest to the issue of paid prioritization..").

Net Neutrality, Unfair Discrimination, & the 2015 Open Internet Order

The primary focus of this paper is the ban on "paid prioritization" promulgated by the FCC in the 2015 Open Internet Order, and the adverse impact it will likely have on the development of new and innovative services that require differential traffic management in order to function properly.[20] However, as context for that discussion, it is useful to first look back at the scholarship and legal history leading up to the Order.

Network Neutrality, Broadband Discrimination and the Four Internet Freedoms

In discussing the merits of differential traffic management on the Internet, it is useful to incorporate the related, but distinct, concept of "network neutrality."[21] That term has taken on significant baggage since it was first coined over a decade ago, and has even gained such prominence as to garner a dictionary definition.[22] The problem with this, and much of the rhetoric in the recent policy debates surrounding net neutrality,[23] is that it bears little resemblance to the idea of net neutrality as first conceptualized back in 2003.[24]

Lexicographers have, understandably, defined "net neutrality" in the most literal sense of those terms: "The idea, principle, or requirement that [ISPs] should or must treat all Internet data as the same regardless of its kind, source, or destination."[25] But that definition does not accurately reflect what Tim Wu was arguing for.[26] Wu was advocating not for a policy that

---

[20] 2015 OIO, ¶ 18.

[21] *See, e.g.*, Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 JOURNAL OF TELECOMMUNICATIONS AND HIGH TECHNOLOGY L. 141 (2003), *available at* http://goo.gl/fDalUr.

[22] *See, e.g.*, *Net Neutrality*, MERRIAM-WEBSTER UNABRIDGED DICTIONARY (last visited Aug. 31, 2016), *available at* http://goo.gl/AZQtCY ("The idea, principle, or requirement that Internet service providers should or must treat all Internet data as the same regardless of its kind, source, or destination."); Mario Trujillo, *Merriam-Webster Adds "Net Neutrality,"* THE HILL (May 27, 2015), http://goo.gl/eOoZq2.

[23] *See, e.g.*, The Oatmeal, *Blog: Net Neutrality*, THE OATMEAL (Nov. 10, 2014), *available at* http://goo.gl/ZMzIR4 ("Net Neutrality . . . simply classifies the internet in such a way that all data, regardless of origin, must be treated equally.").

[24] *See supra* note 21, at 144 (describing net neutrality as "shorthand for a system of belief about innovation policy").

[25] MERRIAM-WEBSTER, *supra* note 22.

[26] Wu, *supra* note 21, at 167–73.

forced ISPs to treat all Internet traffic the same, regardless of what type traffic it is.[27] Rather, he was proposing an unfair competition regime for the Internet, which would tolerate certain forms of ISP behavior (e.g., prioritizing VoIP traffic over email traffic or file transfers) while punishing others (e.g., prioritizing VoIP traffic from one provider over VoIP traffic from another, in exchange for payment) that have the effect of distorting competition and harming consumers.[28]

Now, that sounds like it would be within the purview of the Federal Trade Commission, which uses its authority under Section 5 of the FTC Act to protect consumers, in almost every market, from unfair methods of competition and unfair or deceptive acts or practices.[29] However, it also concerns the Internet, which — comprising a network of networks all communicating with each other via wire and/or radio — is within the realm of the Federal Communications Commission.[30] And by 2004, the Chairman of the FCC (Michael Powell) was willing to embrace a set of net neutrality principles that largely adhered to Wu's vision,[31] with the full Commission formally adopting those principles in a policy statement the following year.[32] It also had the first opportunity to police Internet openness when a small ISP in North Carolina had been blocking the use of VoIP applications, but the company in that case agreed to cease its discriminatory behavior in a

---

[27] *Id.* at 170 ("As the analysis above recognized, the concept of a total ban on network discrimination is counterproductive. Rather, we need distinguish between forbidden grounds of discrimination, those that distort secondary markets, and permissible grounds, those necessary to network administration and harm to the network.").

[28] *See id.*

[29] *See* 15 U.S.C. § 45(a) (giving the FTC authority to police "[u]nfair methods of competition in or affecting commerce, and unfair or deceptive acts or practices in or affecting commerce," by all "persons, partnerships, or corporations, except [*inter alia*] common carriers").

[30] *See* 47 U.S.C. § 152(a) ("The provisions of [the Communications Act of 1934] shall apply to all interstate and foreign communications by wire or radio[.]").

[31] *See* Michael K. Powell, Chairman, Fed. Comms. Comm'n, Remarks at the Silicon Flatirons Symposium on: The Digital Broadband Migration: Toward a Regulatory Regime for the Internet Age, *Preserving Internet Freedom: Guiding Principles for the Industry* (Feb. 8, 2004), *available at* https://goo.gl/xlWVRl.

[32] *See* FCC News Release: FCC Adopts Policy Statement, *New Principles Preserve and Promote the Open and Interconnected Nature of Public Internet* (Aug. 5, 2005), *available at* https://goo.gl/SG7Dj9.

consent decree.[33] It was not until 2007 when an ISP's discriminatory behavior resulted in a legal challenge that actually tested the basis for net neutrality.

Comcast, *Verizon*, and *USTA:* The FCC's Quest for Legal Authority

After tests from the Associated Press and the Electronic Frontier Foundation demonstrated that an ISP (Comcast) had been discriminating against a certain peer-to-peer application (BitTorrent) by surreptitiously throttling its upstream traffic,[34] advocacy groups Free Press and Public Knowledge filed a formal complaint with the FCC calling for it to take action.[35] Comcast's behavior would likely have generated an FTC investigation and potentially a Section 5 violation, but the FCC — then led by Kevin Martin, Michael Powell's successor — decided to assert itself by acting upon the complaint and bringing an enforcement action, despite having only a policy statement and a weak claim of ancillary jurisdiction to point to for authority.[36]

Eventually, the FCC lost in court,[37] and Comcast and BitTorrent settled their dispute out of court,[38] so there was no need (or chance) for the FTC to weigh in. But, by this point, the FCC had gotten a taste for net neutrality, so it set about trying to craft a new order with enforceable rules based on solid legal authority. It first tried using Section 706 of the Telecommunications Act of 1996[39] as authority for rules on transparency, blocking, and unreasonable discrimination,[40] but that effort was largely overturned in court as well, with

---

[33] In re Madison River Communications, LLC and Affiliated Companies, Order, File No. EB-05-IH-0110 (Mar. 3, 2005), *available at* https://goo.gl/0V6lD5; Declan McCullagh, *Telco Agrees to Stop Blocking VoIP Calls*, CNET (Mar. 3, 2005), *available at* http://goo.gl/JGqUxf.

[34] Seth Schoen, *EFF Tests Agree with AP: Comcast is Forging Packets to Interfere with User Traffic*, ELECTRONIC FRONTIER FOUND (Oct. 19, 2007), *available at* https://goo.gl/VdCshP.

[35] Public Knowledge, *Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation For Secretly Degrading Peer-to-Peer Applications* (Nov. 1, 2007), *available at* https://goo.gl/Pr7bEi.

[36] Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications, *Memorandum Opinion and Order*, File No. EB-08-IH-1518, 23 FCC Rcd 13028 (Aug. 1, 2008), *available at* https://goo.gl/CNjSMu.

[37] Comcast Corp. v. FCC, 600 F.3d 642 (D.C. Cir. 2010).

[38] Jacqui Cheng, *Comcast Settles P2P Throttling Class-Action for $16 Million*, ARS TECHNICA (Dec. 22, 2009), *available at* http://goo.gl/lMXjGe.

[39] Telecommunications Act of 1996, Pub. L. No. 104-104, § 706, 110 Stat. 153 (Feb. 8, 1996), as modified by the Broadband Data Improvement Act, Pub. L. No. 110-385, 122 Stat. 4096 (Oct. 10, 2008) (codified at 47 U.S.C. § 1302).

[40] Preserving the Open Internet, *Report and Order*, GN Docket No. 09-191, 25 FCC Rcd 17905 (Dec. 21, 2010) ["2010 OIO"], *available at* https://goo.gl/Vc8ImS.

only the rule on ISP transparency surviving.[41] Then, as the saying goes, the third time was the charm: the FCC finally promulgated a net neutrality order that survived judicial review where both of its predecessors had fallen.[42] For better or worse, in its quest to craft an enforceable net neutrality regime, the FCC reclassified ISPs as common carriers under Title II of the Communications Act, stripping the FTC of its jurisdiction over their broadband Internet access services,[43] leaving the FCC as the sole regulator in charge of policing unfair competition and discrimination by ISPs.

Dissecting the 2015 Open Internet Order

When it was all said and done, the 2015 Open Internet Order did four things: (1) reclassify fixed and mobile broadband as common-carrier "telecommunications services" under Title II of the Communications Act;[44] (2) enhance the 2010 Open Internet Order's transparency rule;[45] (3) adopt a flexible standard to assess unfair competition and discriminatory behavior on a case-by-case basis using a multi-factor test;[46] and (4) adopt bright-line rules prohibiting blocking, throttling, and paid prioritization.[47]

The first of these actions, as alluded to earlier, is questionable, as it takes one seasoned cop off the beat in this area, but the rationale behind it is plain to see.[48] Updating and enhancing the transparency rule is also sensible, particularly after having removed the FTC's ability to police unfair or deceptive advertising by ISPs. Similarly, while one might quibble with some parts of the proposed multi-factor test, the flexible standard the FCC plans to use to police unfair competition and discrimination on a case-by-case basis makes some good sense, as that closely mirrors the FTC's approach to this area and will be best able to adapt to changing circumstances and evolve over time as technology advances, users adapt, and the business practices of ISPs and edge providers change.

What is more difficult to understand is the FCC's choice to supplement the transparency rule and flexible standard with three "bright line" rules specifically prohibiting certain

---

[41] Verizon v. FCC, 740 F.3d 623, 659 (D.C. Cir. 2014).

[42] U.S. Telecomm. Ass'n v. FCC, No. 15-1063 (D.C. Cir. 2016), *available at* https://goo.gl/Wt3T7q.

[43] 2015 OIO, ¶¶ 37–40, 306–433.

[44] *Id.*

[45] *Id.* ¶¶ 23–24, 154–85.

[46] *Id.* ¶¶ 20–22, 133–53.

[47] *Id.* ¶¶ 14–19, 110–32.

[48] Using Title II gave the FCC a stronger base of authority and also, by stripping away the FTC's jurisdiction, avoided the potential problem of overlapping authority and inconsistent regulations from the two agencies.

practices: blocking,[49] throttling,[50] and paid prioritization.[51] The rule on throttling includes an exception for reasonable network management,[52] and it is difficult to see what work it does independent of the prohibitions on blocking and unfair discrimination, but it is not the most objectionable of the three rules. The blocking rule also includes an exception for reasonable network management,[53] and applies only to lawful content, but it still raises difficult Constitutional issues,[54] and arguably conflicts with the Communications Decency Act.[55] However, the focus of this paper is on the third rule, which banned so-called "paid prioritization."

In paragraphs 125–32 of the 2015 Open Internet Order, the FCC laid out its case for why "paid prioritization" — defined as "the management of a broadband provider's network to directly or indirectly favor some traffic over other traffic, including through use of techniques such as traffic shaping, prioritization, resource reservation, or other forms of preferential traffic management, either (a) in exchange for consideration (monetary or otherwise) from a third party, or (b) to benefit an affiliated entity" — should be banned, ex ante, as a per se unreasonable practice.[56]

To justify this proscription, the FCC pointed to support in the record from commenters asserting that differential traffic management is "inherently a zero-sum practice," which will inevitably lead to a bifurcated Internet wherein certain edge providers are relegated to a

[49] 2015 OIO, 30 FCC Rcd 5884 (Appendix A) (codified at 47 C.F.R. § 8.5).

[50] *Id.* (codified at 47 C.F.R. § 8.7).

[51] *Id.* (codified at 47 C.F.R. § 8.9).

[52] 47 C.F.R. § 8.7 ("A person engaged in the provision of broadband Internet access service, insofar as such person is engaged, shall not impair or degrade lawful Internet traffic on the basis of Internet content, application, or service, or use of a non-harmful device, *subject to reasonable network management.*") (emphasis added).

[53] 47 C.F.R. § 8.5 ("A person engaged in the provision of broadband Internet access service, insofar as such person is engaged, shall not block lawful content, applications, services, or non-harmful devices, *subject to reasonable network management.*") (emphasis added).

[54] *See, e.g.*, Geoffrey A. Manne et al., *A Conflict of Visions: How the "21st Century First Amendment" Violates the Constitution's First Amendment*, 13 FIRST AMENDMENT L. REV. 101 (Winter 2015), *available at* http://goo.gl/s7oJqL.

[55] *See* 47 U.S.C. § 230 (giving ISPs authority to engage in "Good Samaritan" blocking of content that either the ISP or the user "considers to be obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable, whether or not such material is constitutionally protected").

[56] 2015 OIO, ¶¶ 125–32.

"slow lane" with degraded performance.[57] The FCC also cited to multiple opposing comments in the record that argued traffic prioritization is not a zero-sum game,[58] but, evidently, did not find these arguments persuasive.

Many feel this decision was unwise, unsupported by substantial evidence, or even contrary to the evidence that was in the record,[59] but, of course, as the expert agency in the field the FCC received *Chevron* deference for its decision and it was upheld by the D.C. Circuit.[60] This paper will now reexamine that particular issue, and argue for the merits of differential traffic management, with the hope that the FCC and other policymakers take it into account down the line when evaluating the practices of ISPs.

On the Relationship Between QoS & QoE: Why Differential Traffic Management on the Internet is Not a Zero-Sum Practice

This paper posits that — contrary to the arguments accepted by the FCC — differential traffic management on the Internet is not a zero-sum practice. Specifically, this paper recognizes the relationship between QoS for edge services and QoE for users, and observes that while these two variables are positively correlated for all edge services, the strength of such correlation differs significantly. For edge services that are particularly sensitive to latency, packet loss, or bandwidth constraints (e.g., live video, online multiplayer gaming), the users' QoE will have a strong direct correlation to the services' QoS. Conversely, for edge services that are comparatively insensitive to latency, packet loss, and bandwidth constraints (e.g., email, software updates), the direct correlation between QoS and QoE will be significantly weaker.

Thus, if an ISP were to prioritize services that fall into the former category, and de-prioritize services that fall into the latter category, the benefits yielded to the former (i.e., higher user

[57] *Id.* ¶ 126 n.287 (citing Mozilla Comments, *supra* note 15, at 20).

[58] *Id.* ¶ 126 n.287 (citing *Protecting and Promoting the Open Internet*, Reply Comments of ADTRAN, GN Docket No. 14-28, at ii (July 15, 2014) ["ADTRAN Comments"], *available at* https://goo.gl/knqUEi and Letter from Justin (Gus) Hurwitz, Assistant Professor, University of Nebraska College of Law, to Marlene H. Dortch, Secretary, FCC, GN Docket No. 14-28, at 1 (Nov. 3, 2014) ["Hurwitz Letter"], *available at* https://goo.gl/0V7kAe).

[59] *See, e.g.*, U.S. Telecomm. Ass'n v. FCC, *supra* note 42, at 25–51 (Williams, J., concurring in part and dissenting in part) ("In fact, as we'll see, the Commission's hypothesis that paid prioritization has deleterious effects seems not to rest on any evidence or analysis."); *see also* Brief for Intervenors for Petitioners TechFreedom et al., U.S. Telecomm. Ass'n v. FCC, No. 15-1063 (Aug. 6, 2015), *available at* http://goo.gl/2nBHDE.

[60] U.S. Telecomm. Ass'n v. FCC, *supra* note 42, at 115.

QoE) would outweigh the harms done to the latter (i.e., slightly lower user QoE), so the network as a whole would work better and aggregate consumer welfare would be increased. For example, in such a system, users streaming live video would be much happier, while users checking their email would be only a little worse off, and may see no significant reduction in QoE. This suggests that the FCC should perhaps reconsider its view of such traffic management practices going forward, particularly as new services with greater QoS demands (e.g., VR-based services) begin to enter the market.

To explain this point, this section first compares different traffic management techniques, before turning to case studies of specific applications that require certain types of traffic management in order to function, and finally assessing the effects these traffic management techniques will have on users through the lens of game theory.

Best Efforts vs. Differential Traffic Management

By default, all ISPs manage traffic as best as they reasonably can.[61] When routers and switches have enough bandwidth to transmit all the incoming traffic as fast as it comes in, everything goes swimmingly. And for those times when networks experience local congestion — with more traffic coming in than can be passed along at any given point in time[62] — they tend to drop packets, at random, knowing that the dropped packets can (usually, but not always) simply be sent again so that the complete data packages can be reassembled at their ultimate destinations whenever all the packets finally arrive.[63]

---

[61] *See, e.g.*, FCC Open Internet Advisory Committee, *2013 Annual Report*, at 77 (Aug. 20, 2013) ["2013 Open Internet Advisory Committee Report"] ("The Internet provides 'best effort' delivery of packets — no guarantees of delivery or delivery time of packets, no guarantees one packet will have the same path/fate as the next.").

[62] Timescales in the realm of computer science are often difficult to conceptualize, as things tend to happen more quickly than in the analog realm. For the sake of simplicity, the timescale this paper defaults to is one second. So, in this instance, there theoretically is some congestion in the network over the course of that second, so packets must occasionally wait in queues for further processing and perhaps even get dropped during transmission.

[63] For example, if some packets from a software update or other file transfer get dropped along the way, they can simply be sent again, compiled at their ultimate destination, and the user likely would not notice any difference (other than the total transfer taking longer than it would have if none of the packets had been dropped during transit), whereas if some packets from a live video stream (e.g., an interactive video-conferencing session, or the live broadcast of a news or sporting event) are dropped during transit, the user (if paying attention) would notice a reduction in quality (e.g., lag, jitter, or lower resolution video), but the missing packets would not be requested to be sent again, because the streaming session would continue on in real time and there would be no use for the packets dropped in the past.

This style of traffic management works great for things like email, file transfers, and, to a lesser extent, basic web browsing. For these services, if QoS suffers, and a few Internet packets get dropped in congestion along the way or arrive all out of order, the user experience is largely the same.[64] In those cases, the emails, files, or web pages would arrive a bit later,[65] but otherwise be unaffected, and the user may not perceive any difference. Indeed, this was a virtue of the early-stage Internet, as it was downwardly-neutral (using the familiar OSI conceptual model),[66] allowing basic Internet services to function effectively despite significant changes in the underlying infrastructure (e.g., the experience of using email on a slow DSL connection and a fast fiber, cable, or wireless connection is largely the same).[67]

However, other Internet services, especially real-time services like telephony and multiplayer online gaming, require a certain level of QoS in order to function effectively,[68]

---

[64] *See, e.g.*, 2015 BITAG Report, *supra* note 10, at 5 ("Customers of Internet access services use those services for a broad range of applications. However, customers rarely notice the underlying transfer of data across the network that enables these activities, except when a performance issue causes a *perceptible* reduction in quality in the application they are using.") (emphasis added).

[65] That is, later than they hypothetically would have if the transmission had been 100% perfect, with no packets dropped along the way and no significant congestion experienced on any of the links carrying the traffic. In reality, congestion — which produces latency and packet loss — is quite common, although it is often sporadic, lasting mere milliseconds. *See, e.g.*, *Inquiry Concerning the Deployment of Advanced Telecommunications Capability to All Americans in a Reasonable and Timely Fashion, and Possible Steps to Accelerate Such Deployment Pursuant to Section 706 of the Telecommunications Act of 1996, as Amended by the Broadband Data Improvement Act*, 2016 Broadband Progress Report, GN Docket No. 15-191, ¶¶ 106–09 (Jan. 28, 2016) ["2016 Broadband Progress Report"], *available at* https://goo.gl/SOQz6T; 2015 BITAG Report, *supra* note 10, at iii ("When TCP transfers a large file, such as video content or a large web page, it practically guarantees that it will create recurring momentary congestion at some point in its network path. This effect exists by design, and it cannot necessarily be eliminated by increasing capacity.").

[66] *See, e.g.*, Microsoft, *The OSI's Seven Layers Defined and Functions Explained* (last visited Aug. 31, 2016), *available at* https://goo.gl/P4lvxN.

[67] *See, e.g.*, Wu, *supra* note 21, at 149; 2013 Open Internet Advisory Committee Report, *supra* note 61, at 77 ("This approach has meant that the Internet is resilient overall, no participating network imposes performance requirements on another, and interconnection between networks is simplified with minimal agreements and commitments required between providers.").

[68] *See, e.g.*, Int'l Telecomm. Union, ITU-T, Telecomm. Standardization Sector of ITU, *Series G: Transmission Systems & Media, Digital Systems & Networks, International Telephone Connections & Circuits; General Recommendations on the Transmission Quality for an Entire International Telephone Connection; G.114: One-Way Transmission Time*, at 1 (May 7, 2003)

and users' QoE degrades significantly as QoS suffers.[69] A best-efforts approach to Internet traffic management — literal net neutrality — is therefore not upwardly neutral, insofar as it discriminates in favor of services that are robust to latency and other QoS disruptions, and against services that are sensitive to latency and require a high level of QoS to work.[70] But real-time, latency-sensitive services have still been developed for the Internet, to varying degrees of success.

Challenges of Using Voice, Video, & Other Real-Time Services

Consider the case of Jeff Pulver and the other entrepreneurs who, building on the work of some brilliant Israeli engineers, pioneered the VoIP services that have now become common on the Internet.[71] VoIP, or IP telephony, requires a certain level of QoS in order to

---

["ITU G.114"], *available at* https://goo.gl/UF0APF ("This Recommendation provides guidance on the effect of end-to-end one-way delay (sometimes termed latency), and an upper bound on one-way network delay. . . . Highly interactive tasks (e.g., some speech, video conferencing and interactive data applications) may be affected by delays below 100 ms[.]"); 2013 Open Internet Advisory Committee Report, *supra* note 61, at 77 ("Increasingly, [over-the-top, edge] services support high-performance hardware on the client end as well as the server end, with attendant expectations of network connections that support their activity. They include applications with particular performance expectations — subject to reduced quality in the face of jitter or high latency, or even any form of timing disruption. A case in point is massive multiplayer action video games, where network-induced delays not only cause deterioration in the video quality experience, but can also get a player killed in a game.").

[69] *See, e.g.*, *id*. at 3, Figure 1 (showing a hypothetical curve representing users' satisfactions with VoIP service decreasing precipitously as latency increases); Marchese, *supra* note 18, at 2–3 (citing Davide Adami et al., *An Applied Research Study for the Provision of a QoS-Oriented Environment for Voice and Video Services Over Satellite Networks*, 25 COMPUTER COMM. J. 1113, 1117–22 (2002), *available at* http://goo.gl/En2KGM) (showing users' reported mean opinion scores of various Internet services to be significantly higher with QoS management than without); 2015 BITAG Report, *supra* note 10, at 6 (citing Markus Fiedler et al., *A Generic Quantitative Relationship Between Quality of Experience and Quality of Service*, 24 IEEE NETWORK 36, 38 (Mar.–Apr. 2010), *available at* https://goo.gl/b9JImo) (showing a generalized shape for the relationship between QoS impairment and the QoE for an application).

[70] Wu, *supra* note 21, at 149–50 ("IP's neutrality is actually a tradeoff between upward (application) and downward (connection) neutrality. If it is upward, or application neutrality that consumers care about, principles of downward neutrality may be a necessary sacrifice.").

[71] *See, e.g.*, ISRAEL21c Staff, *Israel Lauded for Pioneering Work in VoIP*, ISRAEL21C (June 19, 2005), *available at* http://goo.gl/rwxZsJ.

function.[72] In fact, the International Telecommunication Union has long had standards for just what levels of QoS are needed for VoIP,[73] and major networking providers have long had similar internal standards for their services.[74] They recognized early on that, "[f]or VoIP to be a realistic replacement for standard public switched telephone network (PSTN) telephony services, customers need to receive the same quality of voice transmission they receive with basic telephone services — meaning consistently high-quality voice transmissions."[75]

Today, most ISPs offer a standalone VoIP service in addition to general broadband Internet access service,[76] and there are a variety of edge providers offering comparable VoIP services over-the-top ("OTT").[77] These services, while usually crystal clear when a user is calling someone within the same network, often suffer degraded QoS — and, correspondingly,

---

[72] *See, e.g.*, Nick Kephart, *Troubleshooting VoIP QoS and DSCP*, THOUSANDEYES (Oct. 6, 2015), *available at* https://goo.gl/AuXXtp ("Because VoIP traffic is highly susceptible to congestion, loss, latency and jitter across a network, it is a frequent target of complaints and a big time suck for troubleshooting.").

[73] Int'l Telecomm. Union, ITU-T, Telecomm. Standardization Sector of ITU, *Series G: Transmission Systems & Media, Digital Systems & Networks, International Telephone Connections & Circuits — General Recommendations on the Transmission Quality for an Entire International Telephone Connection*, G.114: One-Way Transmission Time (May 7, 2003), *available at* https://goo.gl/UF0APF; Int'l Telecomm. Union, ITU-T, Telecomm. Standardization Sector of ITU, *Series G: Transmission Systems & Media, Digital Systems & Networks; Digital Terminal Equipments — Coding of Voice and Audio Signals; G.729: Coding of Speech at 8 kbit/s Using Conjugate-Structure Algebraic-Code-Excited Linear Prediction* (June 2012), *available at* https://goo.gl/mu1gDa.

[74] *See, e.g.*, Cisco, *Quality of Service for Voice over IP* (June 30, 2001), *available at* http://goo.gl/Twnhnm. Interconnected VoIP and other such services (e.g., VoLTE) typically utilize the integrated services (IntServ) protocols to reserve resources (e.g., bandwidth) and manage QoS so as to deliver high user QoE, rather than the differentiated services (DiffServ) protocols utilized by OTT edge services. These services are allowable under the 2015 OIO, if structured as Non-BIAS Data Services, but may still be subject to FCC enforcement actions under certain conditions. *See* 2015 OIO, *supra* note 12.

[75] Cisco, *supra* note 74, at 2.

[76] *See, e.g.*, Comcast XFINITY, *What is Voice Over Internet Protocol?* (last visited Aug. 31, 2016), https://goo.gl/cyRlcE ("XFINITY Voice service uses VoIP technology to provide you a number of enhanced new features without sacrificing any of your current phone features or the call clarity you expect.").

[77] *See, e.g.*, Michael Muchmore, *The Best VoIP Phone Services for 2016*, PC MAG (Mar. 8, 2016), *available at* http://goo.gl/coq1m0.

lower user QoE — when calls must traverse two or more networks,[78] "[b]ecause each operator sets their own QoS policies[.]"[79] Yet, despite these QoS challenges, OTT VoIP services have become increasingly prevalent in recent years, with many edge providers now integrating VoIP functionality into their existing services,[80] and OTT VoIP services growing tremendously in popularity,[81] especially among younger users.[82]

Similarly, OTT video services have proliferated in recent years, with subscribership and viewing increasing in droves,[83] while subscribership for traditional multichannel video programming delivery ("MVPD") services continues to fall.[84] However, while OTT video-on-demand services that serve cached content (e.g., Netflix, Hulu) can deliver a high user QoE — comparable to that of linear MVPD services — without ISPs implementing QoS differentiation,[85] those that serve live content (e.g., WatchESPN, NBC Sports Live Extra)

---

[78] *See, e.g.*, Kephart, *supra* note 72 (using network path visualization techniques to show how VoIP QoS may degrade when traversing across a network that marks VoIP traffic as a lower-priority DSCP service class than on the originating network).

[79] *Id.*; *accord* 2015 BITAG Report, *supra* note 10, at i ("Notably, traffic differentiation [using the DSCP field] has not been implemented in multi-provider environments, although it is extensively used within specific networks.").

[80] *See, e.g.*, Josh Constine, *Slack Voice Calling Finally Arrives on Desktop*, TechCrunch (Mar. 2, 2016), *available at* https://goo.gl/xrRXCi.

[81] *See, e.g.*, Napier Lopez, *WhatsApp Users Are Making 100 Million Calls Every Day*, THE NEXT WEB (June 23, 2016), *available at* http://goo.gl/IpU18m.

[82] *See, e.g.*, John Fetto, *The $19 Billion Question: Who Uses WhatsApp and Why Are They So Important to Facebook?* EXPERIAN (Feb. 21, 2014), *available at* http://goo.gl/vp5bmU ("While adults of all ages use WhatsApp, users are 15 percent more likely to be between the ages of 18 and 34.").

[83] *See, e.g.*, Nathan McAlone, *Netflix Has More Subscribers Than Its Biggest Cable Competitors Combined, But That Might Be A Bad Thing*, BUSINESS INSIDER (Aug. 2, 2016), *available at* http://goo.gl/NCnr0Q.

[84] *See, e.g.*, Mike Farrell, *Kagan: Telco Losses Drive Pay TV to Record Quarterly Decline*, MULTICHANNEL (Aug. 29, 2016), *available at* http://goo.gl/ctOH7I.

[85] Such services typically utilize a series of techniques, including compression technologies (to minimize bandwidth needs), variable bitrates (to adaptively increase or decrease throughput according to network conditions), buffers (to locally cache content ahead-of-time on a user's device) and/or content delivery networks ("CDNs"), either their own make or from a third party (e.g., Akamai), in order to ensure a good streaming experience and high user QoE even in the presence of QoS disruptions. *See, e.g.*, 2015 BITAG Report, *supra* note 10, at 25 n.7.

are less able to withstand QoS disruptions while still providing high user QoE.[86] With significant caching, such live video services can deliver a seamless viewing experience even in the face of minor QoS disruptions, but such caching can still significantly reduce user QoE when the time delay leads to spoilers (from text messages, social media posts, etc.), which are particularly problematic in the cases of sporting contests and other breaking news alerts.[87]

And there are also other OTT edge services that function in real-time and cannot take advantage of caching. Services like video-conferencing, telemedicine, supervisory control and data acquisition ("SCADA"), certain forms of online multiplayer gaming, and probable forthcoming services like autonomous cars and virtual-reality ("VR") chatrooms, have no ability to use caching to improve user QoE in the face of QoS disruptions because, with those services, content must be pushed from end to end as fast as it is being produced (as with telephony) in order for the services to function properly.[88]

For these services, each QoS disruption (e.g., higher latency or increased packet loss resulting from temporary congestion) may be keenly felt by the user, and QoE will diminish accordingly.[89] In mathematical terms, QoS and QoE are directly correlated for all applications, but the relationship between QoS and QoE is significantly stronger for real-time, highly-interactive services (e.g., VoIP, live video streaming), than it is for other applications that are less interactive and/or less time-sensitive (e.g., web browsing, file

---

[86] *See, e.g.*, Lauren Goode & Sean O'Kane, *Here's the Tech NBC Built to Stream the Olympics — Now Can It Replace TV?* THE VERGE (Aug. 26, 2016), *available at* http://goo.gl/KSGMuk ("To say [NBC] took home the gold in streaming video would also be an overstatement: online viewership was up, but the online experience wasn't quite ready to replace the traditional TV experience.").

[87] Admittedly, users can take actions to avoid such spoilers when watching live video on significant delay, such as turning off their devices or implementing other tools, but these additional steps impose costs on users and remove them from the social experience of watching and reacting to events along with other viewers online, thereby reducing QoE. *See, e.g.*, Fatima Wahab, *Block TV Show, Movie, and Sports Spoilers On Facebook & Twitter [Chrome]*, Addictive Tips (Apr. 26, 2016), *available at* http://goo.gl/vueZx5.

[88] *See, e.g.*, 2013 Open Internet Advisory Committee Report, *supra* note 61, at 80 ("Establishing prioritization of traffic at the access ISP is only going to solve part of the performance problem. Non-interactive services can couple access priority with heavy (and heavily distributed) caching, but that is not applicable in the case of massively multiplayer games. Such OTT services need to have solid network performance between all nodes involved in the interaction, including any transit links.").

[89] *See, e.g.*, Fiedler, *supra* note 69, at 39–40 (comparing the effect of QoS disruptions on QoE for VoIP vs. web browsing, and finding QoS to have a greater impact on VoIP QoE than on web browsing QoE).

transfers).[90] The policy implications of this simple observation for the future of differential traffic management are considered next.

Game Theory & Zero-Sum Thinking

An analytical lens often used in economics (and computer science)[91] is that of game theory, which, in basic terms, is "the analysis of a situation involving conflicting interests (as in business or military strategy) in terms of gains and losses among opposing players."[92] Classic examples of game theory include the famous prisoners' dilemma[93] and Nash equilibrium,[94] but game theory has now permeated popular culture and political debates to the point where it is commonly bandied about to justify policy positions that bear little, if any, relationship to the underlying theory.

The most common example of fallacious reasoning masquerading as game theory involves the concept of a zero-sum game. Accurately stated, a game can be deemed "zero-sum" when, and only when, the gains obtained by one player in the game are offset by equal and opposing losses to another player.[95] However, claims that something is or is not a "zero-sum

---

[90] *See, e.g.*, *id.* (finding the correlation coefficient (R) between QoS and QoE for VoIP to be either 0.998 (Figure 3) or 0.993 (Figure 4), while the correlation coefficient between QoS and QoE for web browsing to be only 0.966 or 0.954 (Figure 5)).

[91] *See, e.g.*, Larry Hardesty, *Gaming the System*, MIT TECH. REV. (June 18, 2013), *available at* https://goo.gl/HCOcES (chronicling the history of game theory and its recent applications to the fields of engineering and computer science). Indeed, given that game theory requires the consideration of systems governed by rules, and ISPs' networks are increasingly managed according to rules defined by software, the fit between game theory and computer science seems quite good. *See, e.g.*, Coevolve, *The Rise of Software Defined Networking*, COEVOLVE (Jan. 21, 2015), *available at* http://goo.gl/geqehz.

[92] *Game Theory*, MERRIAM-WEBSTER DICTIONARY (last visited Aug. 31, 2016), *available at* http://goo.gl/J5OAy5.

[93] *See, e.g.*, Hardest, *supra* note 91.

[94] *Id.*

[95] *See, e.g.*, *Zero-Sum*, MERRIAM-WEBSTER DICTIONARY (last visited Aug. 31, 2016), *available at* http://goo.gl/i5JrV5 (defining "zero-sum" as "of, relating to, or being a situation (such as a game or relationship) in which a gain for one side entails a corresponding loss for the other side).

game" have recently come to pervade political debates,[96] which are typically animated by fears about inequality and unfair competitive.[97]

Thus, it is perhaps unsurprising that commenters like Mozilla suggested to the FCC that paid prioritization must be banned, *ex ante*, because it is "inherently a zero-sum practice."[98] It is undoubtedly true that ISPs could use differential traffic management techniques to discriminate in favor of one edge provider's service and against that of another — say, in exchange for payment or some other sort of kickback — and, in so doing, harm the open nature of the Internet, just as Tim Wu originally forewarned, by warping competition among edge providers.[99] However, leaping from the premise that "whenever a higher priority packet is bumped up in a queue and effectively given priority, every packet that it passes by is left worse off and suffers degraded performance, in the form of higher latency, increased risk of packet loss, or in aggregate, lower bandwidth[,]"[100] to the conclusion that "[p]rioritization is inherently a zero-sum practice,"[101] is a logical non sequitur.

For the reasons laid out above, even if the QoS benefits afforded to one edge service (e.g., more bandwidth, lower latency, or less packet loss) are directly offset by corresponding QoS harms rendered to another edge service (i.e., less bandwidth, higher latency, or more packet loss),[102] that still would not make such discrimination a zero-sum practice, due to the

---

[96] *See, e.g.*, Greg Ip, *The Rise of Zero-Sum Economics*, WALL STREET J. (July 20, 2016), *available at* http://goo.gl/LHQQWZ.

[97] *See, e.g.*, *id.* ("[Trump] sees [trade and the economy] not as a form of cooperation where everyone wins but a contest where someone must lose for someone to win. 'We already have a trade war, and we're losing badly,' he declared last month.

It's not just Mr. Trump who has embraced economics as a bleak zero-sum game; so have Democratic activists. Their platform this year calls the economy "rigged" in favor of the 1%, at the expense of everyone else.").

[98] *See* Mozilla Comments, *supra* note 15, at 20.

[99] *See* Wu, *supra* note 21, at 146 ("A communications network like the Internet can be seen as a platform for a competition among application developers. Email, the web, and streaming applications are in a battle for the attention and interest of end-users. It is therefore important that the platform be neutral to ensure the competition remains meritocratic.").

[100] Mozilla Comments, *supra* note 15, at 20.

[101] *Id.*

[102] And, as other commentators in the Open Internet proceeding argued, even this claim is suspect, for various reasons. *See, e.g.*, ADTRAN Comments, *supra* note 58, at ii ("Sandvine is wrong in claiming that [traffic differentiation] is a 'zero sum game' and that paid prioritization will lead to congestion for others, because Sandvine ignores the fact that ISPs' capacity is not static."); Hurwitz Letter, *supra* note 58, at 1–2 ("There are many reasons that prioritization is not 'zero sum' — indeed, there are circumstances under which prioritization

differing strengths of correlations between QoS and QoE for different applications. Indeed, it is widely agreed among industry experts that certain forms of differential traffic management can benefit some services with minimal or no corresponding harm to other services.[103]

Thus, in light of the observed relationship between QoS and QoE, it appears that differential traffic management, rather than being a zero-sum practice, is actually a positive-sum practice. If an ISP were to prioritize services that are very sensitive to QoS, and de-prioritize services that are less sensitive, the benefits yielded to the former (i.e., higher user QoE) would outweigh the harms done to the latter (i.e., slightly lower user QoE), so the network as a whole would work better and aggregate consumer welfare would be increased. For example, in such a system, users streaming live video would be much happier, while users checking their email would be only a little worse off, and may see no significant reduction in QoE.

Conclusion: Policy Implications & Potential Trade-Offs Going Forward

The foregoing analysis demonstrates that, in light of the relationship between QoS and QoE, differential traffic management on the Internet is not a zero-sum practice. Of course, it is possible that allowing edge providers and ISPs to adopt service level agreements ("SLAs") guaranteeing priority treatment for certain services could produce the sort of negative

---

of some traffic may *improve* the performance of all other traffic. Several reasons for this were presented, including opportunistic use of bandwidth otherwise underutilized by standard TCP congestion control algorithms; avoiding losses and retransmission of sustained, bursty traffic; right-sizing of congestion windows to avoid packet loss; improved retransmission and recovery performance. In each of these cases, there are circumstances under which prioritization of some traffic will reduce congestion faced by other traffic and therefore improve performance of that other traffic.") (emphasis in original).

[103] *See, e.g.*, 2015 BITAG Report, *supra* note 10, at iii ("When differentiated treatment is applied with an awareness of the requirements for different types of traffic, it becomes possible to create a benefit without an offsetting loss."); *id* at 6 ("In many cases, differentiating between traffic flows can improve the QoE for some applications without materially degrading the QoE for other applications."); Fiedler, *supra* note 69, at 40 ("An exemplary user is completely satisfied [with a web browsing experience] if the session time is around half a second. If the data is [sic] delivered faster than that, the user is not able to value this better service quality. Thinking of network provisioning, this opens the possibility to save resources, as it is not necessary to provide better QoS for maintaining the same QoE. This illustrates the potential impact of QoE and the paradigm change in telecommunications networks accompanied by the consideration of QoE instead of QoS.").

externalities that would justify a continued ban from the FCC,[104] but it is also possible that such arrangements could produce the sort of positive externalities that should be championed by the FCC,[105] and as of yet the agency can point to no economic studies that suggest either outcome is more likely than the other, or that the benefits of the paid prioritization ban exceed its costs.[106]

It seems that, in the end, the FCC was swayed by populist rhetoric and fearmongering into adopting the precautionary principle, and banning paid prioritization agreements across the board, subject to very limited exceptions,[107] out of an overabundance of caution and fear of the unknown. The FCC claims that it "do[es] not seek to disrupt the legitimate benefits that may accrue to edge providers that have invested in enhancing the delivery of their services to end users[,]"[108] but its ban on paid prioritization does just that: it disrupts the ability of edge providers and ISPs to deliver high quality services to end-users.

Of course, this is not the end of the matter. The FCC's ban on paid prioritization is the law of the land for now,[109] but scholarship and research in this area march inexorably forward,

---

[104] For example, if an ISP were allowed to accept payment (above and beyond typical interconnection fees) in exchange for providing priority treatment for some services, it may have less incentive to invest in improving the quality of its treatment for lower-priority services, or it may raise barriers to entry for potential competing services. *See, e.g.*, 2013 Open Internet Advisory Committee Report, *supra* note 61, at 66 ("[If allowed to offer specialized services], broadband providers might constrict or fail to continue expanding network capacity allocated to broadband Internet access service in order to provide relatively more capacity for specialized services.").

[105] For example, if an ISP were allowed to accept payment (above and beyond typical interconnection fees) in exchange for providing priority treatment for some services, it may have added incentive to invest in overall network capacity, which would improve the level and quality of service available to all forms of Internet traffic, including lower priority traffic, as it all utilizes the same IP infrastructure.

[106] *See, e.g.*, U.S. Telecomm. Ass'n v. FCC, *supra* note 42, at 41 (Williams, J., concurring in part and dissenting in part) ("One prominent critic of the ban on paid prioritization — Timothy Brennan, the Commission's chief economist at the time the Order was initially in production, who called the rules 'an economics-free zone' — offered an alternative [proposal] that addressed these concerns [about negative externalities].") (citing L. Gordon Crovitz, *"Economics-Free" Obamanet*, WALL STREET J. (Jan. 31, 2016), *available at* http://goo.gl/PxJh7g).

[107] 2015 OIO ¶ 130 ("The Commission may waive the ban on paid prioritization only if the petitioner demonstrates that the practice would provide some significant public interest benefit and would not harm the open nature of the Internet.").

[108] 2015 OIO ¶ 128.

[109] It is worth noting that, while the FCC won the initial appeal of its order before the D.C. Circuit panel, as of the time of this writing, that litigation is still ongoing, and the D.C.

continually building upon the ideas and work of those who have gone before. Indeed, the bright minds at the Internet Engineering Task Force ("IETF") have been contemplating potential uses for differential traffic management since at least 2006,[110] and their work continues to this day,[111] despite the FCC's proscription. Furthermore, if the agency suddenly acquires some scruples with regard to changing its mind and reversing previous policy positions,[112] it need not bend over backwards to reconsider its ban on paid prioritization.

In the wake of the 2010 Open Internet Order, the Specialized Services Working Group within the FCC's Open Internet Advisory Committee did a "forward-looking case study" considering "the challenge of supporting applications that have a requirement for enhanced service qualities that cannot today be met over the Internet."[113] Among the potential examples of differential traffic management on the Internet that could meet the challenge of providing high QoS to services that require it in order to deliver high user QoE, the Committee considered an "[o]pen-standards based approach to signaling requests and requirements throughout the network[.]"[114]

The Committee suggested that, "[a] future approach might be to ensure that there are open standards and best practices that are developed to support highly interactive traffic in

---

Circuit sitting en banc and/or the Supreme Court may yet overturn some or all of the 2015 Open Internet Order on procedural, statutory, or Constitutional grounds.

[110] *See, e.g.*, J. Babiarz et al., *Request for Comments 4594: Configuration Guidelines for DiffServ Service Classes*, THE INTERNET SOC'Y (Aug. 2006), *available at* https://goo.gl/32ONt8.

[111] *See, e.g.*, 2015 BITAG Report, *supra* note 10, at 4 ("While they have been used in specific networks . . . the IPv4 ToS field and the IPv4/IPv6 DSCP, have not been deployed or used across network interconnects for both engineering and economic reasons, and would require the harmonization and cooperation of the relevant network operators. Proposals to that end are being discussed in the IETF, however.") (internal citations omitted); R. Geib & D. Black, *DiffServ-Interconnection Classes and Practice: Draft-IETF-TSVWG-DiffServ-Intercon-9*, IETF TRUST (Aug. 28, 2016), *available at* https://goo.gl/U2lttG.

[112] *See generally* U.S. Telecomm. Ass'n v. FCC, *supra* note 42, at 1–25 (Williams, J., concurring in part and dissenting in part) (arguing that the FCC's reclassification of broadband Internet access service from a Title I, information service, to a Title II, telecommunications service, under the Communications Act of 1934, was unlawful because it failed to properly account for the reliance interests reasonably engendered by the previous classification); *see also* FCC v. Fox Television Stations, Inc., 556 U.S. 502, 515 (2009) (explaining that sometimes an agency faces a higher burden in justifying a policy position, such as when "its new policy rests upon factual findings that contradict those which underlay its prior policy; or when its prior policy has engendered serious reliance interests that must be taken into account").

[113] 2013 Open Internet Advisory Committee Report, *supra* note 61, at 77.

[114] *Id.*

general, and perhaps some level of mutually-cooperative signaling of performance preferences that works across network domain boundaries in the Internet."[115] This also is very similar to what the Broadband Internet Technical Advisory Group ("BITAG") proposed in its uniform agreement report last October: "BITAG continues to recommend transparency when it comes to the practices used to implement the differential treatment of Internet traffic."[116]

Such an approach — mandating transparent disclosures about ISP traffic management practices and allowing consumers to make their own choices based on those disclosures and their individual preferences — would be far preferable to the *ex-ante* ban the FCC adopted on paid prioritization in the 2015 Open Internet Order. To again quote from the FCC's Open Internet Advisory Committee, "In all of this, perhaps the most important thing for the FCC to consider is the distinction between challenges and solutions for today, versus opportunities tomorrow."[117] A policy that bans certain practices outright, without proof that the benefits of such a ban will exceed its costs, is unwise and antithetical to the history and spirit of innovation and entrepreneurship that have made the Internet so tremendously successful.

Rather than using the precautionary principle and adopting a "Mother, may I?" regime for monitoring particular forms of business arrangements and Internet traffic management, the FCC should continue allowing the Internet to evolve and change with the times, as technologies advance and consumer preferences adapt, and not stultify it with burdensome and unnecessary regulation. A simple transparency requirement, coupled with a case-by-case approach to assess particular issues as they arise and guard against the sorts of discriminatory behavior and unfair competition "net neutrality" was originally intended to prevent,[118] would be adequate to ensure the future of an open Internet, while also being more conducive to investment and innovation throughout the Internet ecosystem. Hopefully the FCC will reconsider its view of traffic management practices, remove the ban on paid prioritization, and allow innovative new edge services (e.g., virtual reality chat rooms, or other highly interactive services that must traverse multiple hops in real-time with high QoS in order to deliver high user QoE) to flourish.

---

[115] *Id.*

[116] 2015 BITAG Report, *supra* note 10, at iv, 29 ("Specifically with respect to consumer-facing services such as mass-market Internet access, network operators should disclose the use of traffic differentiation practices that impact an end user's Internet access service. The disclosure should be readily accessible to the public (e.g., via a webpage) and describe the practice with its impact to end users and expected benefits in terms meaningful to end users.").

[117] 2013 Open Internet Advisory Committee Report, *supra* note 61, at 81.

[118] *See* Wu, *supra* note 21.