

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

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

Comments of

David Clark, William Lehr, and Steve Bauer¹
Massachusetts Institute of Technology

¹ David Clark, ddc@csail.mit.edu; William Lehr, wlehr@mit.edu; and Steve Bauer, bauer@mit.edu. We are submitting these comments as individuals. The opinions expressed herein are those of the authors alone.

Table of Contents

1. Introduction.....	5
2. Framing our point of view	7
2.1. Recognize all the relevant actors	7
2.2. The objectives of the proposed intervention.....	8
3. Innovation on top of the Internet platform—applications and services.....	9
3.1. Utility	9
3.1.1. The status of QoS.....	10
3.1.2. Over-provisioning	10
3.1.3. Other forms of service differentiation	11
3.2. Stability	11
3.2.1. The role of the ISP	12
3.3. Business risk	13
3.3.1. Offering a basic Internet access service.....	14
3.4. The layered nature of applications and content	14
4. Fair access to content	15
4.1. Content modification in the net	15
5. The health of the Internet.....	17
5.1. The growing cost of usage	17
5.2. Business arrangements between ISPs and owners of content.....	18
5.3. Abuse is a two-way street	18
6. Fear of abuse	20
6.1. Discrimination today.....	20
6.2. Other examples of discrimination.....	21
6.3. The need for a rule	22
7. Cost-effective un-hindered access to the Internet experience.....	22
7.1. The “dirt road” fear	23
7.2. A philosophy of control	23
8. Constructing workable rules	24
8.1. Build on other venues for debate	24
9. Summing up and some final thoughts on potential rules.....	25

Executive Summary

We are submitting these comments in response to the Notice of Proposed Rulemaking issued by the Federal Communications Commission (FCC) on October 22, 2009, in the matter of *Preserving the Open Internet and Broadband Industry Practices*.²

In the enclosed, we offer our thoughts on particular aspects of the FCC's NPRM as individuals with substantial relevant experience in the technical, business, and policy environment that has helped foster and shape the Internet. We are not lawyers and our intention is not to offer precise wording for regulatory rules, but rather to comment more generally regarding those aspects of the NPRM which we find more or less promising.

Specifically, we approve of the FCC's effort to establish clear regulatory guidance regarding its intention to promote and protect an open and competitive Internet. We believe it is appropriate that the FCC's efforts in this regard be based on general guidelines rather than overly detailed specifications because the Internet continues to evolve. We believe that being too specific at this point would likely result in too much rigidity in the policy framework, rendering it less able to provide helpful direction and more likely to interfere with the efficient working of the market process. We view the FCC's reliance on general rules with the potential for *ex post* enforcement and further refinement or clarification on a case-by-case basis as consistent with the transition toward light-handed, market-based regulation and as an appropriate strategy for promoting a healthy Internet ecosystem.

In the attached comments, we identify some of our concerns with the six proposed draft rules which codify the four Internet policy principles originally adopted in 2005³ and add a fifth *non-discrimination* rule and a sixth *transparency* rule. We believe the first three principles which focus on protecting the rights of end-users to access the content, use the applications, and connect the devices of their choice (subject to these being lawful and consistent with reasonable network management) and the fourth principle which focused on protecting competitive choice among different components of the Internet value chain (i.e., network providers, application and service providers, and content providers) are appropriately focused on important concerns. In the four years since these were originally promulgated, it seems that most stakeholders have embraced these principles as worthwhile and important. This consensus is valuable and often hard to achieve in policy debates, and should be counted as a success.

In its efforts to transform the original four principles into rules to govern the behavior of network service providers, we believe the FCC has focused too narrowly on end-users and network

² See *Notice of Proposed Rulemaking*, In the Matter of Preserving the Open Internet and Broadband Industry Practices, Federal Communications Commission, GN Docket No. 09-191 and WC Docket No. 07-52, Released October 22, 2009 (hereafter, "NPRM"). The opinions expressed herein are those of the authors alone.

³ See NPRM, paragraphs 5.

service providers – or Internet Service Providers (ISPs). We believe the rules need to be interpreted not as obligations applicable only to ISPs but as general statements of the sort of market outcomes the FCC will seek to protect – namely, competitive choice. ISPs are not the only market actors whose behavior might pose a threat for an open and healthy Internet.

The FCC should anticipate that in the future, the service platform provided by the ISP (the so-called “Internet” service) is not the only platform over which there will be debates about openness and discrimination. Any statements of objectives should be done in a way that recognized this fact, and takes it into account in the intended way

With regards to the newly proposed non-discrimination rule, we do not believe this is appropriate at this time. It is premature to attempt to specify appropriate network management practices for ISPs. The technical and business communities are in the process of evolving new technologies, standards, and business practices to manage the rapidly evolving broadband Internet. ISPs are expected to play an important role in traffic management in a healthy Internet. While bad practices may be employed and cautious policy oversight is warranted, we believe that the market is generally better suited for guiding the determination of what constitute acceptable practices at this time.

With regards to the newly proposed transparency rule, we see merit in trying to improve the quality of information available in the marketplace regarding the details of broadband service offerings; however, we recognize that ISPs may regard the specifics of how they manage their networks to be strategically sensitive and proprietary information. We recommend limiting required disclosure to network management techniques that alter the semantics of Internet transport.

1. Introduction

This document is submitted to the FCC in response to the Notice of Proposed Rulemaking “In the matter of Preserving the Open Internet” (FCC 09-93).⁴

While we agree with the FCC that the objective of an open Internet is important, we have some reservations about the specific approach taken by the FCC to achieve this goal, and some of the assumptions and reasoning that underpin that approach. We offer these comments as individuals with significant experience in the technical, business, and policy environment that has shaped the Internet. We are not lawyers and we take no position on whether the original Internet policy principles⁵ need to be made into “rules” in order for the FCC to act. We are concerned with the set of behaviors that are deemed to be acceptable or unacceptable as specific rules are formulated.

First, the current proposed rules are stated in terms of obligations on Internet access service providers, or equivalently, Internet Service Providers (ISPs). Stating specific obligations has the potential short-term benefit of clarity. However, we feel that this approach requires a clear articulation of the high-level objectives that underpin those obligations, in a form that the current NPRM does not achieve. Lacking a clear articulation of objectives, over time the obligations can be mistaken for the objectives, and take on a life of their own at a time when the obligations need to evolve and adapt to the changing technical and business context.

Second, the rules, as stated, raise many uncertainties and ambiguities in our minds that should be resolved to the extent possible before they are made final. We find the proposed rules too constraining; they impose limits on ISPs that are beyond what is justified to preclude unacceptable behavior. We agree that there have been examples of inappropriate and unacceptable behavior—for example blocking certain content and applications in use on the Internet without the permission of the relevant parties. We agree that it is appropriate to define such behavior as inappropriate. On the other hand, we do not see an epidemic of such behavior. We believe that overly aggressive attempts to provide ex ante limits on ISP network management practices may prevent useful and beneficial behavior and, paradoxically, may be anti-innovation

⁴ See *Notice of Proposed Rulemaking*, In the Matter of Preserving the Open Internet and Broadband Industry Practices, Federal Communications Commission, GN Docket No. 09-191 and WC Docket No. 07-52, Released October 22, 2009 (hereafter, "NPRM"). The opinions expressed herein are those of the authors alone.

⁵ Appropriate Framework for Broadband Access to the Internet over Wireline Facilities; Review of Regulatory Requirements for Incumbent LEC Broadband Telecommunications Services; Computer III Further Remand Proceedings: Bell Operating Company Provision of Enhanced Services; 1998 Biennial Regulatory Review – Review of Computer III and ONA Safeguards and Requirements; Inquiry Concerning High-Speed Access to the Internet Over Cable and Other Facilities Internet Over Cable Declaratory Ruling; Appropriate Regulatory Treatment for Broadband Access to the Internet Over Cable Facilities, Policy Statement, 20 FCC Rcd 14986, 14987–88, para. 4 (2005) (Internet Policy Statement).

and actually interfere with the evolution of a healthy and open Internet. We have attempted to suggest alternative formulations that steer clear of that possibility.

Third, the proposed rules do not explicitly recognize and catalog all of the relevant actors and their interests. The ultimate goal is to promote an open and healthy Internet which allows as much competitive choice as possible at each point in the Internet value chain. The original Internet policy principles were focused on protecting the consumer and were framed as “consumer entitlements.” The consumer is not the only relevant actors whose interests need to be protected and whose behavior might threaten the health or openness of the Internet. Application developers, content providers, equipment manufacturers, and the ISPs also have interests and potential behaviors that should be considered in the formulation of appropriate rules. Taking into account the interests of all these actors, we believe that there will be few absolute prohibitions on one or another that will result. The rulemaking must reflect a balancing act.

For this reason, we agree with the approach of defining general rules and then relying on *ex post* enforcement and refinement/clarification on a case-by-case basis; this approach is preferable to attempting to define overly specific *ex ante* rules regarding what constitutes acceptable behavior. This is more consistent with the evolution of policy toward light-handed, market-based regulation, in which greater reliance is placed on market processes to determine the details of market outcomes. This approach is consistent with good regulatory practice in light of the complexity of the Internet value chain and the fact that the Internet and the industries that create it continue to evolve. While it may be possible to define specific behavior that is clearly acceptable or unacceptable, there will be a large middle ground, and the FCC policy must not attach a presumption of misbehavior to this region—this region is where much innovation will occur and where market forces need scope to act with a degree of freedom from regulatory constraint.

Based on our concerns, we offer a point of view that leads to an alternative framing of some of the proposed FCC rules (related to open access for applications and for content), and an argument that one of the rules (rule 5, related to discrimination) is not necessary to achieve the desired outcome. We find that rule 5 is unlikely to be effective in limiting undesirable discrimination, but may limit many beneficial activities. It imposes overly strong *ex ante* constraints on ISP traffic management practices. Rule 5 appears to us to be a first step in regulating interconnection agreements among ISPs and between ISPs and content providers, and we believe that the FCC should not be seeking to regulate these arrangements at this time and as part of this proceeding.

Our analysis of innovation in applications and services leads to the conclusion that the rights of the various parties, both consumer and application designer, would be best served if ISPs are permitted to offer a range of services that may enhance or impair certain applications, as long as those ISPs offer at least one service that provides access without hindrance to any application of the user’s choice.⁶ We analyze the relationship between the ISP and the application designer, and

⁶ Specialized or small ISPs may be exempt from the requirement to provide at least one tier of broadband Internet access service that allows the user to access the content, use the applications, or connect the device of their choice without interference or discrimination by the ISP. Later we address the concern that this basic broadband service might be of such poor quality as to be effectively useless.

note that while the current regulatory posture seems to presume a specific concern -- namely, that ISPs are in a position to hinder applications -- abuse can occur in both directions. Applications can abuse the norms of usage that are important to the economics of sharing and that are traditionally associated with the Internet service. We argue that each party should be subject to obligations and policing. Appropriate rules should establish the framework for allowing the FCC to act if the desired market outcome of a healthy, open, and competitive Internet is threatened, regardless of the identity of the agent of that threat.

Our analysis of access to content leads to the conclusion that the interesting case is not where a willing sender and a willing receiver agree to transfer lawful content (in which case the ISP should not interfere), but in the case where the interests of the sender and the receiver are only partially aligned. We find that the rule proposed by the FCC does not capture this critical and common situation.

In summary, we believe the FCC has an important role to play in providing regulatory oversight and guidance to the Internet marketplace, and that this is best executed through general principles or "rules" that are subject to *ex post* enforcement, clarification, or refinement on a case-by-case basis. This approach leaves significant scope and flexibility to market actors and processes to determine the range of acceptable network management behaviors. We explain why we believe that the best means to achieve the desired objectives is to find ways to encourage cooperative and collective behavior among the actors, rather than to focus on the exceptions where unacceptable activities may have occurred.

The rest of our comments are organized into seven sections. Section 2 explains why we believe a multi-stakeholder perspective is important to formulate appropriate policies for promoting a healthy and open Internet. In Sections 3 and 4 we focus on the roles and concerns of application developers and content providers, respectively. In Section 5, we consider the challenges confronting ISPs and the need to protect and promote incentives for adequate infrastructure investment. Section 6 focuses on the concern over inappropriate discriminatory practices by an ISP, while Section 7 addresses the "dirt road" fear that allowing ISPs to offer differentiated, and potentially discriminatory, services in addition to a basic tier of Internet service might result in the unacceptable deterioration of the basic broadband service tier. We offer concluding remarks in Section 8.

2. Framing our point of view

2.1. Recognize all the relevant actors

The FCC's original Internet policy statements from 2005 focused on the consumer, but the consumer is not the only actor that is directly influenced by potential rulemaking.⁷ Today we must be concerned with the consumer, but also the developers of new applications and services,

⁷ See NPRM, paragraph 101: "Although the question of Internet openness at the Commission has traditionally focused on providers of broadband Internet access service, we seek comment on the pros and cons of phrasing one or more of the Internet openness principles as obligations of other entities, in addition to providers of broadband Internet access service."

creators of content, *and* the ISPs. In the future, we may see new actors emerge, and new relationships among those actors. When we consider potential rulemaking, the policy should specifically recognize the diversity of these actors, and not try to cast the policy using one of these, for example the consumer, as a proxy for the others. Nor should it use obligations just applied to one of the actors as a way to balance their interests.

Any analysis of objectives will reveal that while these actors have interests that are sometimes aligned, they are also at times at odds. Any defensible articulation of objectives should explicitly reflect a balancing of the interests of all the actors. Using one as a proxy for the others obscures this analysis. Therefore, we believe that the FCC should consider including language in the rules that makes it clearer that the obligation to avoid behaviors that might threaten the health, openness, or competitiveness of the Internet include all the key stakeholders, and are not limited to the ISPs.

2.2. The objectives of the proposed intervention

The original policy statements of the FCC, now being refined as rules, were stated as entitlements of the consumer. The current framing of the rules is not stated in terms of an objective, such as protecting the consumer, but in terms of specific obligations to be imposed on ISPs. Being clear about obligations is important, but unless that obligation is placed within the context of a framing objective, there is a risk that over time the obligation is confused with the objective, and becomes the focus of attention in and of itself. Over time, we can expect both the Internet and the structure of the industries that create the Internet to evolve and morph. The obligations should be viewed as malleable; appropriately framed objectives should be the foundation for ongoing reasoning about the obligations.

The NPRM identifies four high-level issues.⁸ In the following sub-sections, we address these after slightly rearranging them to better tease apart the interests of the different actors, as follows:

- Preservation of vigorous innovation on top of the Internet platform.
- Fair access to content.
- The continuing health of the Internet itself.
- Cost-effective un-hindered access to the Internet experience.

Before proceeding with this discussion, however, it is worth offering a few thoughts on the "future adequacy of competition and market forces." While we believe that the FCC appropriately identifies the promotion of effective competition as a worthwhile objective, we do not believe that the basis for justifying the promulgation of the rules proposed in the NPRM requires a finding that the ISPs or other relevant stakeholders have (or are likely to have in the future) market power that poses a threat to competition. We recognize that concerns over market power, or equivalently, a lack of adequate competition are legitimate and likely to remain so. However, this is not the only potential market failure that might threaten the health of the

⁸ See NPRM, paragraph 61: "The arguments in this area have largely revolved around four issues: (1) how best to promote investment and innovation; (2) the current and future adequacy of competition and market forces; (3) how best to promote speech and civic participation; and (4) the practical significance of network congestion to the other considerations.

Internet and may call for active policy interventions. Whether ISPs do indeed have market power and in which markets, or what other actors may have market power and in which markets are interesting questions that are appropriately beyond the scope of this proceeding. While market power findings may influence the specific remedies employed to address the potential consequential problems, our recommendations regarding the rules considered in this NPRM are consciously agnostic about the adequacy of competition.

3. Innovation on top of the Internet platform—applications and services

The platform quality of the Internet—its general utility for designers of a broad range of applications and higher-level services, has been a driver for a great deal of innovation and economic benefit. This fact argues that there is social benefit in preserving the desirable qualities of that platform. However, it is important to carefully consider what aspects of that platform make it suitable for innovation. The simple concept of “neutral” does not capture what the desirable features actually are. We argue that from the perspective of innovators that develop new applications and services, the important features of the Internet platform include the following:

- Utility and generality
- Stability over time and consistency across regions
- Predictable business risk

Put differently, innovators at the higher levels may be deterred because the platform is not as useful as it could be, because it changes beneath them in ways that requires them to respond in costly ways, as well as the fear that they are unexpectedly denied the use of the platform.

3.1. Utility

The Internet's value ultimately resides in the usefulness, or utility, it provides to end-users. The Internet needs to offer an appropriate quality of service (QoS) experience to end-users. We are pleased that the FCC has recognized that different network circumstances (e.g., congestion) and applications (e.g., delay sensitive and insensitive) may benefit from differentiated services.⁹ Some neutrality arguments seem to imply that a totally neutral platform with a single, best-effort service is best for stimulating innovation, but we believe that providing different service qualities for different sorts of applications is important. On the one hand, we hear from providers of games and VoIP that their applications would work much better if they could be assured a service with bounded delays, and on the other hand, we see one provider of peer-to-peer services, BitTorrent, designing their application to intentionally defer its own traffic in favor of games and VoIP. According to their press releases, they do this in order to make their product more appealing to the market.¹⁰ We see BitTorrent's decision to voluntarily modify the behavior of a

⁹ See, for example, the discussion in the NPRM at paragraph 137.

¹⁰ In a press release dated September 15th, 2008, BitTorrent stated the following:

“BitTorrent DNA is able to automatically detect game play on a user’s system and through BitTorrent’s advanced proprietary congestion control technology eliminates any impact on game

sophisticated application to allow it to defer to other traffic as evidence of the success of the market process and as direct evidence against arguments that network-based tools for quality differentiation are anti-innovation.

3.1.1. The status of QoS

Today, the Internet protocols include standardized mechanisms (e.g., RSVP, DiffServe) to provide quality of service (QoS). Many applications would seem to benefit from these enhancements to the customary platform. In private networks based on the Internet protocols (corporate intranets and the like) these tools are used, seemingly to good effect. But they are not offered as part of the customary service platform of the public Internet.

If one asks why ISPs do not offer these services on the public Internet today, one gets a range of related answers: they cannot devise a business model and a set of interconnection agreements that allow them to charge for the service end-to-end, and they worry that they will be accused of violating the current norm of neutrality on the public Internet. Regulators should encourage, rather than hinder, any sort of business negotiation that would allow ISPs to offer QoS across the public Internet, which would include negotiation of interconnection agreements (and payment for those agreements) for the carriage of enhanced services.

3.1.2. Over-provisioning

Some observers seem to argue that a preferred alternative to adding QoS is simply to expand capacity, or equivalently, over-provisioning of the network so that congestion does not occur. The corollary of this is that regulations that impede an ISPs ability to implement QoS mechanisms may cause limited harm. We believe that this line of reasoning is flawed—it is based on an over-simple understanding of how the Internet works. Since TCP tries to go as fast as possible unless it is being artificially throttled (as does occur today in some cases), congestion will occur somewhere along the path, if only in the server itself.

ISPs in the U.S. today offer access services with capped sending and receiving rates (with the caps typically between a few and a few tens of megabits per second). A properly configured PC today can sustain transfers at a rate of a few hundreds of megabits per second. So any time a user transfers a large file, the network software on the PC will attempt to execute the transfer at a rate that will exceed the speed cap set by the ISP, and will trigger a transient form of congestion sometimes called “self-congestion”. This congestion need have nothing to do with whether other users are also transferring data at the same time, and this sort of congestion is not a signal that “something is wrong”; it is a natural consequence of the Internet operating the way it was designed to operate, since (under most circumstances) the network protocols of the Internet (specifically TCP) try to go as fast as possible.

play as well as other sensitive real-time applications including Voice over IP (VoIP) calls, streaming media, or Internet browsing.”

BitTorrent is using a protocol they call uTP, which is based on their DNA algorithm, to enable BitTorrent traffic to temporarily back-off when it anticipates that doing so will alleviate network congestion.

Self-congestion is often the cause of some of the “congestion problems” that have been seen in the Internet today (e.g., high variations in latency across access links). For example, if a single user runs two applications at once—a latency-sensitive application such as a computer game and a bulk transfer application, the two may interfere with each other. QoS mechanisms are a way to allow multiple applications running on behalf of a single user to interact in beneficial ways when some of them want controlled latency, and some of them want maximum throughput. Improving the way different applications interact is an enhancement that benefits all parties.

3.1.3. Other forms of service differentiation

In addition to the above, there are a number of other mechanisms for providing applications with the service differentiation they may desire. These include differentiated routing which selectively exploits lower latency or higher capacity paths; replicating traffic over diverse paths for high-availability applications; and selective, link-level retransmission across lossy links depending on the application’s need for loss-free delivery. Experimentation with services such as these should not be precluded out of a fear that they might be misused.

3.2. Stability

The neutrality debate does not usually include any discussions of stability, but it is the nature of a successful platform that there is a tension between stability and innovation. The costs of application development and support are increased if the platform on which an application runs changes unpredictably or too rapidly over time, or behaves differently in different circumstances (public vs. private networks or in different regions of the world.) On the other hand, it is through innovation in the services provided by the platform that the network itself evolves. The interests of different application providers with respect to maintaining the stability of the Internet may not be aligned with each other or with the ISPs. In general, one will see a complex interplay among the various actors through which the platform evolves at a rate that balances their various desires. To date, the Internet has demonstrated its ability to grow in scale significantly and evolve in response to market forces, without significant regulatory interference in how traffic is managed. As a general matter, we prefer to see the market determine the rate and direction in which Internet protocols evolve.

In practical terms, the Internet is not the same everywhere. Most obviously, in different regions its peak speed can vary by five or more orders of magnitude. Latency can also differ widely. These are well-understood variations that application designers accept.

The discussion above concerning differentiated QoS illustrates the balance of issues here. Mechanisms to provide differentiated QoS are present in some networks (many corporate intranets), but not in the public Internet. However, the Internet QoS mechanisms were designed so that the application designer need not redesign the application to deal with their presence or absence. Turning on QoS to achieve reduced variation in latency, for example, just improves the performance of latency-sensitive applications, with no (or minimal) explicit need to redesign the application to accommodate the mechanisms. Further, Internet QoS mechanisms are based on well-documented standards, so any application designer who needs to know how they work can easily find this information. The design of QoS as a service enhancement was undertaken by the

technical community in a way that tried to avoid imposing a major burden on application designers, but which could be exploited when it is available.

3.2.1. The role of the ISP

To understand the balance between stability and platform evolution, one must consider both the needs and expectations of application and service designers, as well as the needs and expectations of ISPs themselves.¹¹ Specifically, from a regulatory point of view, the question is whether regulation should attempt to define what is acceptable and unacceptable innovation with respect to the platform service of the Internet. Our conclusion is that while there are specific forms of innovation that can be deemed acceptable in advance, much experimentation that might occur here will fit into that middle ground where one cannot expect to write specific rules in advance to differentiate what is acceptable and unacceptable.

The examples above suggest two characteristics of an innovation that would support a presumption that the innovation is appropriate for the basic Internet service offered by an ISP.¹² First, if an innovation in the service does not disrupt the semantics of the platform on which the application depends, but just makes it perform in a way more suited to the objectives of the application designer, that should be acceptable. Second, if the innovation is based on standards that are approved by a suitable standards body after suitable deliberation, and are well documented, that sort of innovation also should be acceptable.¹³ What the ISP should be expected to disclose (consistent with the transparency rule) is any aspect of the service innovation that will have an effect on the application. Internal aspects of the implementation of the service (which allow the ISP to do a better job or a more efficient job) should not need to be disclosed, and can form the basis of ISP service differentiation. This rule may offer only very limited opportunities for provider-specific proprietary innovation, but that is the price of our desire to preserve the Internet as a platform well-suited for application innovation.

It is important that this rule be stated in terms of its outcome, rather than in terms of specific prohibitions. For example, it might seem that deleting a packet in transit by an ISP could be prohibited, as it seems as if it is a hostile act. However, the snoop module in an architecture proposed by Balakrishnan¹⁴ improves the performance of TCP across lossy wireless networks by deleting packets (specifically duplicate TCP acknowledgements). So any discussion of

¹¹ To be clear, just as some application developers and some ISPs may disagree as to the appropriate pace of platform innovation, different application developers or different ISPs may disagree among themselves, as well.

¹² As noted earlier, so long as the ISP offers an appropriate basic Internet service, it is presumptively free to offer such additional services as it finds appropriate without being limited by the proposed rule which is intended to protect the stability of the platform for application developers.

¹³ Of course, this presumes that the standards organization that approves the standard is not subject to capture.

¹⁴ Balakrishnan, H., Seshan, S., and Katz, R. H. 1995. Improving reliable transport and handoff performance in cellular wireless networks. *Wirel. Netw.* 1, 4 (Dec. 1995), 469-481.

acceptable and unacceptable behavior should be based on outcomes, not on mechanistic obligations and limitations.

Finally, it is worth noting that today, in practical terms, the component that adds the most undocumented variation to the Internet platform is the Network Address Translation device, or NAT. The NAT, an afterthought not well specified by standards, is a substantial burden in cost and complexity for many application designers. It has not prevented innovation, and can be seen as demonstrating that absolute stability of the platform is not a necessary objective of regulation. As we balance the expectations and obligations of the various actors, absolutes in prohibition are not a necessity, but perhaps only a convenience to easy enforcement.

3.3. Business risk

We used the phrase “predictable business risk” to describe one of our desired objectives. We are not suggesting that regulation should have the goal of minimizing business risk. We recognize that innovators may be interested in high-risk experiments if they perceive the potential of high returns. The narrower question that seems to be the center of regulatory debate is whether regulators should intervene when one of the relevant actors behaves in a way that increases the business risk for another of the actors. Non-neutrality of a sort that allows a platform owner unilaterally to block or degrade certain innovations will certainly add additional uncertainty to the risk assessment for innovators. We see an example of this sort of discrimination today with the iPhone, where Apple retains the right to veto applications developed for the iPhone. While some developers have abandoned the iPhone as a platform for innovation, there are many developers who have accepted the risk that Apple may veto their applications. The large number of application developers for the iPhone suggests that discrimination by a platform provider and innovation need not be inconsistent. However, if one concludes that the platform owner has significant market power, then there are reasons to be concerned that the platform owner might engage in socially harmful discriminatory practices to the detriment of competition. So, the concern that discrimination may harm innovation remains legitimate but it is unclear that strong anti-discrimination controls on ISPs are the best way to address this concern.

A more interesting question arises when we try to balance the rights of the consumer, the ISP, and the application designer. We can take as a rough starting point that the consumer should be able to use any application they choose. But if the ISP offers them a bargain in which they agree not to use some set of applications, is that acceptable? In the U.K., T-Mobile sells two forms of wireless broadband access service to hook up a laptop. One, called Mobile Access Plus, costs £15/month and does not permit voice applications like Skype. The other, called Mobile Access Extra, specifically permits Skype. It costs £30/month.¹⁵ These two offerings illustrate how the ISP and the consumer might bargain; it could be described as putting a price on openness. Under what circumstances would such a service offering be acceptable?

One question we might ask of this situation is whether our answer would depend on the relative prices. If the cost of the open version was £5 more, £15 more, or £100 more, would we answer

¹⁵ See <http://www.t-mobile.co.uk/services/mobile-broadband/what-does-it-cost/>, visited 12 December, 2009.

differently? Of course, if we concluded that the market for ISP access were not adequately competitive, we would be more concerned, but as noted earlier, the existence of market power does not provide a suitable *a priori* justification for precluding making such differentiated service offerings available.

As a general matter, we believe this sort of alternative service offering should be acceptable, and that the consumer can be protected from abusive behavior, so long as the ISP also offers a basic Internet access service that is "open" in the sense that it does not constrain a subscriber's choice of applications. It is not essential that this "open service" be the default offering or offering most commonly adopted by broadband subscribers of an ISP. What matters is that is at least one of the options so we can be assured that its absence is not retarding useful innovation.

3.3.1. Offering a basic Internet access service

If an ISP offers a basic Internet access service at a reasonable price, then it can as well offer other services that enhance or degrade the platform: enhance by means such as QoS, or degrade by means such as blocking certain applications. So long as any application restrictions or modifications that affect the semantics of the platform are adequately explained to consumers and application developers, the offering of such additional services should be acceptable.

Of course, this implies the need to define what constitutes an acceptable version of a basic Internet access service and what constitutes reasonable pricing for such a service. We do not intend such an obligation to be burdensome for small or specialty ISPs, but to be most important in situations where there is a substantial basis for presuming that the ISP may have a significant degree of market power.

The NPRM does not attempt to define what constitutes broadband service. Presumably, this will be addressed in part by the FCC's broadband plan. While we recognize the importance of defining what it means to provide a basic Internet access option, we do not undertake commenting on that here as it is beyond the scope of the NPRM.

3.4. The layered nature of applications and content

The problem of defining what constitutes an appropriate version of basic Internet service is complicated by the fact that the platform continues to evolve. For example, the Web is an application that runs on top of the Internet. Facebook is an application that runs on top of the Web. Facebook, in turn, is a platform for third party applications that use its API (application programming interface). Where to draw the line between what constitutes the Internet platform and applications or other platforms that may ride on top of the Internet is non-trivial and evolves over time.

With these thoughts in mind, we believe the FCC should anticipate that the Internet platform and industry structure that supports it and depends on it will continue to evolve but not necessarily in ways that are fully foreseeable today. The generality of the rules and the clear statement of objectives will help the rules remain relevant and robust in the face of changing industry and Internet circumstances.

4. Fair access to content

The second Internet principle we wish to address concerns fair and open access to (legal) content. Again, we believe that this is a market outcome worth preserving. Our goal here is to distinguish among a number of different circumstances that will warrant different treatment.

First, there is a spectrum from what might be called “citizen content” at one end—political speech, freedom of expression and the like, through content that is free on the Internet but commercial in that it is paid for indirectly, perhaps by advertisements, and finally content offered for sale—typified today by premium video channels, commercial music, or first-run movies.

With respect to the content we have characterized as “citizen content”, while today ISPs, as private speakers, can legally discriminate and selectively block this sort of content, we agree with the FCC that it is a valuable social goal that the ISPs not do so. The question, as always, will be with respect to content that is lawful but offensive to many customers of the ISP, who do not want to encounter such content, even by accident.

The proposed FCC rule states that an access provider may not prevent any of its users from sending or receiving the lawful content of the user’s choice. This simple formulation is problematic in that it seems to make sending and receiving into two, unrelated actions. But they are two halves of a single action. Generally, when the interests of the sender and receiver are aligned, and one party wants to receive the lawful content that the other is sending, we believe that the right balance of interests is that the ISP should not interfere.

The more interesting case is when the interests of the two parties are *not* aligned. What if the sender wants to send (or “offer”) the content, but the potential receiver has no desire to receive it. Or perhaps the receiver is uncertain about the content—the receiver would like to receive email but not receive spam or viruses. In this case, how are the interests of the various parties balanced, and what are reasonable roles (obligations and expectations) for the ISP in this case? Again, we see that a reasonable outcome will be a balance of interests, not an absolute rule.

The FCC identified one specific example of this situation: spam. (Note that spam—unsolicited bulk email—is lawful unless it is fraudulent in some way. For the government to make spam illegal would be to violate the first amendment rights of the spammer.) Even though the proposed FCC rule allows a user to send the content of his choice (which would seem to prevent the blocking of lawful spam), the FCC makes clear that ISPs should be able to block spam, presumably with the permission of the receiver, without worries that they have violated the rights of the sender. We are not clear what the general rule is of which this is a specific example.

4.1. Content modification in the net

Here is another specific example of a misalignment between sender and receiver that may provide useful insight in our effort to define an appropriate general rule. Many Web pages cannot be viewed well on the small display of a mobile device, so some mobile operators interpose software that reformats the content on the way from the server to the device. This reformatting

seems to serve a useful purpose, but what sorts of modifications are acceptable¹⁶? One form of modification seems to have nothing to do with mobile devices, and has been quite contentious—the rewriting of URLs in a page to replace the embedded advertising content with ads placed by the ISP. Where is the line between acceptable and unacceptable content modification, and on what basis can a regulator draw this line?

In this particular case, one might get into trouble if one looked only at the rights of the receiver (the consumer) to have the content of his choice. Given a choice, many consumers might want the ads removed altogether. In fact, one could imagine an even more extreme form of intervention (which has not to our knowledge happened) in which the ISP sold a premium service to the consumer in which they stripped all the ads out of all content.

This case illustrates the question of balance we raised earlier. Should an ISP be allowed to carry out any sort of modification of content if the recipient gives permission in their service agreement? Or do the creator of content and the creator of an application have some intrinsic rights that must be preserved? With respect to content, one might argue that these rights are embodied in the creative rights of the content author. But if stripping advertisements is not acceptable, but stripping (potentially legal) spam is, what is the general rule?

We believe that the resolution of this question represents another of these “middle regions” in which *ex ante* rulemaking is problematic. Specific case-based review will be required to determine the correct balance of interests. However, some initial effort to identify and recognize the rights of the content provider, rather than depending on the “subscriber” as a proxy for those other actors, will add clarity to the situation. There is still the tension among these actors, and this tension is fundamental—in the open world of the PC (as opposed to the more closed wireless device world) consumers often download third-party software to strip out ads. We do not believe that consumers should be forced to look at advertisements. But at least, when we recognize the content producer as an actor, we have the tension correctly framed.

One possible policy formulation is that if the consumer does not want to receive something at all, that is the receiver’s choice, but if the receiver does agree to receive something, it should be delivered in the form that the sender generated, or subject only to such modifications (e.g. reformatting to fit a small screen) that the sender accepts. But this formulation still raises questions about removing attachments from mail, deep linking to web sites, and so on. ISPs should not be expected to resolve these sorts of issues on their own.

The FCC should recognize that there needs to be a forum (or in practice more than one) in which the balance of these issues can be debated, and some form of resolution reached which defines acceptable behavior (business practices) for ISPs. It should also be recognized that experimentation by ISPs will be an important part of learning how the various interests in the market should be balanced.

¹⁶ For one proposal for defining acceptable and unacceptable behavior, see *Guidelines for Web Content Transformation Proxies 1.0*, from the W3C consortium, at <http://www.w3.org/TR/ct-guidelines/>.

In the concluding section, we offer some thoughts on how the FCC's proposed rules to reformulate the Internet policy principles focused on ensuring open access to content might be improved.

5. The health of the Internet

This objective is focused on the Internet itself—the world-wide assembly of links, routers, radios, etc and the protocols running on top of them that provide the platform we have been discussing. Society needs continued investment in access facilities (wired and wireless) if a growing share of our population is to have the broadband experience available to them, and if the performance and service quality of the broadband experience is to improve over time to match the improvements we see in other parts of the IT world—the improved performance of processors, the declining cost of storage, and so on.

The construction of last-mile access facilities is capital-intensive.¹⁷ Financing these investments presents a difficult challenge. Funds could in principle come from many sources—the private sector, the government, or from the users themselves. Private sector investment could take many forms—facilities-based competition, unbundling of a pervasive single infrastructure, and so on. Perhaps the highest-level debate in this area is which of these options are viable, and whether regulatory policy should have as its goal picking among these outcomes.

At the present time, the U.S. finds itself in a situation where there is duopoly facilities-based competition in many markets between wired legacy telephone and cable television providers, while some markets (e.g. rural) are underserved or unserved. Fixed and mobile wireless services, whether terrestrial or satellite based, as well as potentially other services (e.g., broadband-over-powerlines) have the potential to offer additional competition in many markets. There have been a number of experiments in public and public-private investment, and the government (through the ARRA stimulus legislation) is adding public-sector money to improve unserved and underserved areas. We do not see that the current FCC rulemaking has the objective of reshaping this palette of approaches. Under this assumption, the health of the private-sector access provider is a critical part of the overall equation. So the balance the regulator must strike in any intervention is to find rules that prevent unacceptable behavior by ISPs while giving them the business opportunities that justify continued investment. In working out this balance, we should consider what obligations, if any, should be placed on other actors.

5.1. The growing cost of usage

With the rapid emergence of Internet video, we see levels of traffic in access networks increasing. These increases in usage imply real costs to ISPs, and we can expect ISPs to respond to these costs. One way is by adding some sort of usage quota to the definition of their service tiers, as well as peak rate caps. We should assume that the introduction of such usage quotas into the market will generate new fears of discrimination and customer abuse, but at the same time, we should recognize that dealing with the cost of usage is a legitimate objective. More

¹⁷ Further, many of the costs are labor costs, which are not subject to any sort of Moore's Law reductions over time.

specifically, usage quotas will introduce new modes of negotiation and bargaining among the different actors—content providers and hosting services, the subscriber, etc.

Moreover, at the same time that traffic-related usage costs are growing, broadband penetration is saturating. This means that the revenues contributed by monthly fixed subscription fees will cease growing, putting pressure on ISPs to seek new opportunities for generating revenue. This may entail a shift in pricing away from flat rate services toward usage-based charging.

5.2. Business arrangements between ISPs and owners of content.

The owners of content are often not a party to an explicit business agreement with the subscriber's ISP, because they contract with some sort of hosting service that in turn arranges the necessary interconnection agreements. However, other sorts of business arrangements will emerge, beyond interconnection agreements. Recently, Comcast and Time Warner, Inc (the provider of content, not the cable company) have announced that they will make some premium video content available to their Internet customers if those customers have first purchased access to that same content through their cable contract.¹⁸

Should we fear this as a signal of dangerous integration between ISPs and content producers, or a constructive example of the integration of the ISP into the content distribution value chain? One can argue that this arrangement brings benefit to the consumer, who can get the same content through multiple channels while paying for it only once. Given that this content is considered "premium" content, the owner is not likely to agree to give it away. Whatever the distribution channel, this is likely to be "paid" content and the ISP may be in good position to participate in implementing an appropriate payment mechanism.

Again, we could ask some hypothetical questions about this arrangement. Should we be concerned if, as part of the agreement, the ISP agreed to give this content preferential delivery treatment? This would seem to be another example of an explicit agreement among the relevant parties to diverge from the basic delivery service. Because of the potential that such an agreement might be beneficial to all parties, we oppose an *ex ante* ban on such contracts.

5.3. Abuse is a two-way street

The focus of the FCC NPRM seems to be to prevent the ISP disrupting the stability of the platform and harming the application/content designer. What about the opposite: when the application designer tries to disrupt the platform to its advantage?

If application designers view the platform as an enemy, they can try to defeat the commonly understood "rules of the road" in order to improve their performance at the expense of other applications. One form of behavior that might be classified in this way is opening up lots of parallel TCP connections and splitting the data to be sent across those several connections. Since

¹⁸ See http://www.usatoday.com/tech/news/2009-12-15-comcast-cable-online-viewing_N.htm.

the normal congestion behavior of the Internet is to limit all TCP flows equally, the user with more flows gets more capacity. At worst, this leads to an arms race of more and more flows¹⁹.

A more serious form of aggression would be to tinker with the tuning parameters of TCP so that it responds more aggressively when it receives signals of congestion. (Or, as an alternative, build an application that uses its own transport protocol that is more aggressive.) This has happened in the past in crude ways: some applications were designed not using TCP, but using the raw datagram interface of UDP, and did not respond to congestion signals at all. This behavior led to public outcry, the designers were convicted in the court of “techie” public opinion, and the IETF called for the imposition of a norm that any transport protocol be “TCP-friendly”, which means it be no more aggressive than the normal “as shipped” TCP.²⁰

It is well understood (if regretfully) that if public opinion and social coercion cannot enforce such a rule, then the ISPs may have to become policemen, and enforce norms of aggression in some way. A growing consensus appears to be emerging in the IETF – the principle standardization body for the Internet – that ISPs need to play a bigger role in traffic management to deal both with issues of abuse and simple usage.²¹ Any proposals for limiting ISP behavior should recognize this shift—if ISPs are prohibited from discrimination based on application, then they will have to impose limits on the user as a whole if the user (perhaps unknowingly) invokes an aggressive application. This leaves the user having to figure out that the reason they are getting poor service overall is that they are running an aggressive application that is causing their overall service to be degraded.

In the tussle between the ISP (as the provider of a platform) and the application (as the utilizer of that platform), we have seen abuses going in both directions. If we take the word *neutrality* as some sort of code-word for “acceptable, non-disruptive behavior”, we need to consider obligations on both parties and define the extent that each have the right to police compliance with industry norms. Regulators should not impose obligations on ISPs that gives ISPs no right to protect themselves from applications and users that abuse the norms of the platform. This last point—that abuse can go both ways—is critical to understanding what is happening today, and help us understand how to draw balanced rules.

We would argue, as we did in the case of dealing with content, that the ISP should be permitted to play the role of policing aggressive applications, but they should not be in the business of

¹⁹ There is a strong limit of diminishing returns: TCP connections tend to provide poor performance if they do not have enough data to send.

²⁰ IETF RFC 3884, “TCP Friendly Rate Control (TFRC): Protocol Specification” <http://www.ietf.org/rfc/rfc3448.txt> for instance states that “a flow is ‘reasonably fair’ if its sending rate is generally within a factor of two of the sending rate of a TCP flow under the same conditions.”

²¹ Some examples of these efforts include <http://www.ietf.org/dyn/wg/charter/alto-charter.html> and <http://trac.tools.ietf.org/area/tsv/trac/wiki/re-ECN>.

unilaterally making the laws they enforce. There may free-rider or coordination problems associated with defining appropriate business practices. The discussion of what constitutes acceptable and unacceptable application behavior should be debated in other venues, just as with ISP behavior.

6. Fear of abuse

The NPRM²² alludes to fears that ISPs may abuse any flexibility they may have to discriminate among traffic. The general form of these fears is that ISPs may discriminate among different application and content providers to favor one over another.

The discrimination could occur at the point where a content provider connects to the ISP, inside the ISP as packets transit the network, or at the access point where the subscriber connects. And the discrimination could take many forms. It could involve direct modification of the packet flow, or variation in pricing in one or another form. The NPRM proposes a specific non-discrimination rule to prevent some forms of such discrimination. We will argue that this rule is not appropriate: it does not eliminate discrimination, and prevents useful forms of behavior.

6.1. Discrimination today

First, we want to consider a particular sub-case where the content provider is directly connected to the ISP in question, and has an interconnection agreement with the ISP.

Consider this simplified interconnection situation: an ISP 1, a consumer A, and two content providers X and Y, all of whom are customers of ISP 1. The fear is that ISP 1 might block or degrade the communication between A and (say) X because ISP 1 has some reason to favor provider Y. In the most extreme version of the story, Y might be a service offered by ISP 1 itself: ISP 1 is both an intermediate carrier of traffic between A and X and a provider of its own competing service Y. But a wide range of business relationships might cause ISP 1 to regard X and Y differently, and to treat them differently.

Before we speculate about what sorts of abuses and remedies might arise here, it is important to look first to the current market, where the Internet today is rich with discrimination of this sort—with respect to interconnection the Internet is anything but neutral. If we look at an ISP (our ISP 1 in the story above) and consider what sort of negotiation arises when other actors (such as our X and Y) request to connect to that ISP, we see a rich space of negotiation and bargaining, and we see a lot of what an economist would call discrimination. The forms of discrimination include discounts for capacity at higher link speeds, and discounts (including payment-free) for so-called *peering* connections, where the content provider connects to the ISP only for the purpose of getting access to the customers of the ISP.²³ Much of this is presumably efficient and so not the sort of potentially harmful discrimination that the FCC is seeking to protect against.

²² See section IV.A.3.b generally, and paragraph 71 specifically.

²³ When ISPs connect, they sometimes choose to *peer*, and peering is sometimes settlement-free or revenue-neutral, and sometime not: one party may pay the other. When major content providers (such as Content Delivery Networks, or CDNs) negotiate to connect to ISPs, the same sort of bargaining occurs. Sometimes CDNs connect to ISP on a revenue-neutral basis, sometimes they pay.

If two providers (X and Y in our example above) bargain to a different outcome with respect to the terms of interconnection, it stands to reason that the customer A may see a different experience when it communicates with X and with Y. The ISP is not doing something secretive and manipulative to the traffic going across the network, it has just implemented a different Service Level Agreement with the two of them. If X has purchased a small connection, the traffic from X to A will flow slowly, even if A has an access connection with high capacity. This situation might be seen as a mismatch of service at the two ends, but this may be what was desired, and the Internet protocols are designed to work in this context, to “do the best they can” given the attributes of the various parts of the path, including the X and the Y connections.

The network neutrality debate has largely not focused on the point of interconnection among ISPs or between ISPs and content providers. However, we might ask what acceptable and unacceptable discrimination might be. X and Y may have different access agreements with different sorts of services. This seems reasonable. If ISP 1 sells to X but refuses to sell to Y an essentially similar service at an essentially similar price, we might ask whether this is acceptable. To evaluate the situation in detail we might need to know all the aspects of the two candidate agreements, all the contextual business concerns and the like²⁴. Today, any attempt to explore this space gets tangled in debates over what might be meant by “essentially similar”, and interconnection agreements have become very complex.

We believe that it would be inappropriate for the FCC to attempt to regulate interconnection agreements today and/or as part of this proceeding.

6.2. Other examples of discrimination

The “fifth rule” of the FCC states that a provider of access service must treat lawful content, applications and services in a nondiscriminatory manner. This wording seems to imply that the only locus of discrimination about which the FCC is concerned is the access link. What happens in the network or at the point where the content provider connects is not governed by this rule. The FCC offers one example: the ISP should not be able to bargain with various content providers to accept a fee in exchange for favoring the traffic of that provider across the access link. We do not believe that this is a useful or balanced prohibition.

First, this rule does not capture the range of mechanisms for discrimination that might arise. For example, we are moving toward a future where the service agreement of the consumer (at the point of access) is defined as much by the usage cap as by the peak rate. As a result, we could easily imagine an arrangement in which a content provider pays an access provider to carry traffic to the subscriber without having that traffic count against the usage quota of the subscriber. This arrangement would not involve providing enhanced or prioritized delivery. It would be a beneficial bargain in many cases for all concerned—providers of high-value, high volume content might be quite prepared to pay a fee to allow the subscriber to receive the

²⁴ A somewhat detailed discussion of the current status of interconnection negotiation in the Internet can be found in an earlier paper, Faratin, Peyman, Clark, David D., Bauer, Steven, Lehr, William, Gilmore, Patrick W. and Berger, Arthur, *The Growing Complexity of Internet Interconnection* (December 31, 2008). *Communications & Strategies*, No. 72, p. 51, 4th Quarter 2008.

information without worries about exceeding a monthly quota. From our point of view, this would be beneficial, rather than unacceptable discrimination.

Second, the FCC rule, as stated, interacts with QoS in a bad way. If we agree that QoS is valuable (as the FCC perhaps implies²⁵) but the provider of the content cannot pay to obtain that service across the access link, then it would seem that unless it is provided for free, the consumer would have to pay to obtain it. In other words, the outcome is a “receiver pays” system, and “receiver pays” systems create security problems in which an attacker can send unwelcome traffic for which the subscriber pays. “Sender pays” schemes are much closer to “self-policing” and require less auxiliary mechanism to make them abuse-proof.

6.3. The need for a rule

The non-discrimination rule seems to have the objective of protecting the content provider and/or application developers. We are not convinced that the potential for abusive discrimination is manifest enough to require a rule. What we see today is legitimate economic discrimination, which is a natural market phenomenon. We feel that many of the specific hypothetical examples are based on an over-simple view of the Internet value chain in the Internet. For example, focusing on content providers, we observe that many small providers of content do not host their own content, but contract with hosting services and content delivery networks (CDNs) to host their content, and this trend seems likely to accelerate with the emergence of large data centers and cloud computing. It is these intermediate providers that need to bargain with ISPs for interconnection arrangements. In this case, there could be two levels of potential discrimination. In one, the ISP may use deep packet inspection to single out one kind of content from another within some hosting service, which would almost certainly violate the service agreement between the ISP and the hosting service. In the other, the ISP may discriminate among the hosting services themselves, which is already happening today without evidence of major barriers or undesirable outcomes.

7. Cost-effective un-hindered access to the Internet experience

This objective is the last of the four with which we started this discussion. This objective tries to capture the needs of the user, as opposed to those of the creator of applications, services and content, and the ISPs themselves.

If the market for Internet access were highly competitive, then the discipline of competition might more easily be accepted as sufficient to protect the user. However, because a significant share of the costs of constructing local access facilities are fixed and do not vary with the level of penetration that an operator realizes in a neighborhood, total costs might be lower with fewer – perhaps a single – facilities-based provider. In the extreme, one might argue that last-mile networks are a natural monopoly, at least in some smaller, less dense markets.

²⁵ Paragraph 106 states that they do not want to prevent subscribers from purchasing “different services”. Paragraph 137 specifically asks for comment on the value of QoS. We have stated our position that explicit QoS is pro-innovation, and should be encouraged, not just permitted.

At the same time, observation of the current access market in the US would suggest that the duopoly we find in parts of the country (cable vs. telco) is a driver of investment in the platform, with increasing performance as a strong component of competition. We see the current shape of the U.S. access market as a strength of our approach, not a weakness.

While some of the justification for intervention by the FCC may be based on the argument that access technology is a bottleneck facility, certain socially desirable outcomes may require regulatory intervention, independent of the level of competition.

7.1. The “dirt road” fear

One of the fears that has been expressed about the Internet is sometimes dubbed the “dirt road” fear—the fear that the ISPs will under invest in the basic Internet service, perhaps at the same time offering enhanced or alternative services at an increased price, with the eventual outcome that consumers are induced to switch to this new service and thus increase the total revenues of the ISPs.

The Internet tools for QoS are sometimes seen as aggravating this fear, since those tools can be used to create new service tiers that, while they benefit selected applications, hinder others, to the point that those other applications become unusable.

One form of protection is to allow the subscriber and content provider, rather than the ISP, to determine which applications get which service quality. If the ISP offers different treatment of packets, but does not get to select which applications get which service, then the ISP gets an enhanced business opportunity, but other actors get the control over it.

If additional protection is required, we would suggest the approach that was already embodied in our proposal that at least one of the service options offered by an ISP that offers broadband service to the general public be a "basic" Internet access service that is suitably undifferentiated. If the content provider has requested that the ISP provide enhanced service across the access link, the consumer should be able to opt out. So we make a distinction between “who pays” and “who controls”.

In the end, with or without tools for service differentiation, ISPs may choose to under-invest if they do not see a return on that investment. We cannot use regulation of the sort proposed here (or by the FCC) as a tool to mandate investment. We must look to the competition in the market, and the expectations that we offer the ISPs to benefit from investment, as encouragement to invest; and, if shortfalls remain, to universal service policies to address those gaps.

7.2. A philosophy of control

Our proposals here, taken generally, have the objective of giving the ISPs a richer role as business partners in the Internet value chain. We recognize that there are some that seem to believe that if regulators do this, it will represent the “camel’s nose in the tent” that lets ISPs exploit their market power. Our alternative point of view is that richer integration into the value chain will lead to partnership feedback that stabilizes their place in the value chain. ISPs will find both new revenue opportunities and increasing business constraints on their ability to act independently in ways that disrupt the interests of other actors.

Through this approach, we believe that policy can protect the subscriber by recognizing the rights of the other actors, rather than the other way around.

This is consistent with the transition to light-handed, market-based regulation. The FCC plays an important role as a referee but most of the detailed decision-making is left to the market. If the market process delivers the appropriate outcomes then the FCC will be like the lonely Maytag repairman in the television advertisements having little to do which would be a good outcome for all concerned.²⁶

8. Constructing workable rules

The vague specification of the service that is provided by the Internet (best effort, few quantitative measures) has been critical in allowing the ISPs to reduce cost by exploiting statistical resource sharing. At the same time, its vague nature has allowed both ISPs and application designers to argue that something they are doing, which seems to bring them benefit at the expense of other actors, is within the bounds of acceptable behavior with respect to this vague specification. We agree with the FCC that it will be complex, constraining and probably ineffective to try to make precise rules separating what behavior is and is not acceptable. We should try, to the extent possible, to define areas that are clearly acceptable and clearly unacceptable, but these *ex ante* rules will of necessity leave a broad grey area. We have to accept this, and allow for a gradual process of *ex post* review to clarify what is and is not acceptable in specific cases. To compensate for this, we should try to state the high-level objectives that will guide this sort of review, as we have tried to do. This approach is desirable and consistent with the transition to market-based, light-handed regulation. It remains a work-in-process that will evolve as the Internet and market evolves.

8.1. Build on other venues for debate

Today, the first line of defense against misbehavior has not been formal regulation or rulemaking, but outcry from the technical community, which does not seem to have problems telling right from wrong in a practical sense. The emergence of the call for “TCP-friendly” behavior in response to certain sort of application-level abuse seems to have worked.

To the extent that behavior (both by ISPs and by application designers) is standards-based, it has at least passed through some court of public opinion. Of course, there are lots of venues that might be willing to issue standards, if this outcome became important. But the invocation of a recognized and accepted standard is one argument that behavior is within the accepted norms.

To the extent that behavior goes beyond what is specified by a standard (which does not mean that it is malicious—standards have never tried to capture all aspects of behavior), it seems as if the court of public opinion is the first place to try behavior.

²⁶ The advertising campaign for Maytag washing machines claimed that the repairman was lonely because Maytag washing machines required so fee repair calls (see http://www.tvacres.com/admascots_maytag.htm).

9. Summing up and some final thoughts on potential rules

The FCC's NPRM seeks to instantiate the four Internet policy principles that were originally promulgated in 2005 as entitlements for Internet end-users to access/use the applications, content, and devices of choice (rules 1, 2, and 3) and to have choice in all three (rule 4) as rules that define obligations on ISPs. Moreover, the NPRM proposes adopting two additional rules: a discrimination rule (rule 5) that imposes restrictions on the traffic management practices and contractual relationships that an ISP may employ; and a transparency rule (rule 6) that requires an ISP to disclose its traffic management practices.

In this document, we have offered comments on this overall approach and these rules, re-ordering and reframing several of them. Our goal is not to opine on the legal necessity or precise language of the rules since we are not lawyers, but rather to highlight concerns we have with the rules as proposed and to suggest areas where we think the rules and the process might be improved.

First, we would like to affirm our view that the FCC should and does play an important role as the leading regulator for the Internet. We recognize the need for a regulator, while at the same time, approving the general trend toward light-handed regulation, which shifts more of the detailed decision-making responsibility to the Invisible Hand of the market. The FCC needs to transition to being a referee – helping to establish the rules of the road and establishing broad objectives – but not trying to define *ex ante* detailed outcomes in the grey areas such as network management and service provisioning where the Internet continues to evolve actively. For these reasons we approve of the NPRM's approach in proposing general rules and relying in an *ex post* review process to provide further clarification and refinement, and to allow the FCC to intervene should abuses occur. In addition, we believe that any rules the FCC approves should consider not just ISPs and end-users, but all of the stakeholders whose behavior might impact the health of the Internet.

Second, with respect to rule 4 (the right to competitive choice), we think this is part of a high-level framing goal: light-handed regulation and markets work best when there is competition. How much competition is viable in each segment of the value chain and in each market is not something that the FCC can pre-ordain, and so we are unclear whether this is really something that needs to be enshrined in a rule. Perhaps it is better to think of this as a goal that helps provide a standard against which to judge FCC actions. That is, the FCC will seek to promote as much competition as is feasible in the Internet and will consider the competitive implications of its actions (or lack of action) in the context of its consistency with that goal.

Third, with respect to discrimination rule (rule 5), we believe that the FCC has failed to make a sufficient case for why such a rule would be beneficial at this time. As proposed, we see the potential for more harm in deterring potentially beneficial innovations such as expanded deployment of QoS mechanisms by ISPs than benefit in addressing the risk that ISPs might engage in harmful discrimination. We do not recommend adopting the discrimination rule.

Fourth, with respect to the transparency rule, we can see benefit in requiring some additional disclosure of how ISPs are managing traffic, especially when failure to offer such disclosure

might plausibly mislead end-users or application developers as to the differences between different Internet access services. We might summarize the key issue for ISP disclosure as follows:

Any service variant that modifies what applications the subscriber can use must be fully disclosed to and accepted by the subscriber. Any modifications that affect the semantics of the platform in ways that are visible to the application designer must be fully documented for the application designer.

Fifth, with respect to the first three rules which are focused on ensuring freedom to use/access/connect the applications/content/devices of choice without hindrance by the ISP, we think these need to be re-framed. Our discussion focuses on applications and content providers, and we do not discuss the issues that may arise with respect to connecting different CPE. Our concern is that the rules as proposed in the NPRM seem overly restrictive. We can imagine ISPs beneficially offering multiple tiers or versions of Internet access services that may vary in the extent to which they deviate from traditional single-class, best-effort, TCP-friendly transport. While we believe there is merit to the Internet and marketplace of preserving the option to use the traditional type of Internet access service (at least today and into the near-term future), we do not think that this is the only type of service that is consistent with promoting effective competition and innovation and the overall health of the Internet.

Consequently, as long as ISPs offer at least one version of a service that supports the traditional type of undifferentiated best-effort Internet access service, they should be allowed to offer other forms of Internet access that may deviate from this (e.g., block certain applications or offer enhanced treatment for certain classes of traffic). We believe this approach will allow the Internet to continue to evolve most effectively and consistent with the changing needs of the marketplace. This approach does not preclude future action by the FCC should circumstances change materially.

The above option implies that the FCC will undertake to provide guidance as to what constitutes an acceptable version of a basic Internet access service and what it means to make such a service available at a reasonable price. We expect that the FCC's broadband plan will offer some further insight into what this might be, and do not view this proceeding as the right place to address what an appropriate specification of what would constitute an acceptable basic Internet access service offering.

As long as such a basic service option is available to end-users and the transparency rule is applied so that users, application developers and content providers understand the relevant traffic management practices that differentiate the basic service offering from what other service offerings the ISP may make available, that that will allow end-users to use the applications they wish.

Finally, because we suspect that the outcome of this process will be a set of rules that are intended to define the range of acceptable ISP behaviors as it relates to mediating the relationships between end-users and the applications they choose to use and content they choose to access and the sponsors/developers of the applications and content, we offer some further

guidance in how those rules might be reformulated to make them more permissive with respect to potentially beneficial forms of innovation. We hope the spirit of the sort of change we think would be beneficial is captured in the following possible reformulation of rule 1 regarding the ISPs possible role in influencing the choice of applications:

Any provider that offers “Internet access service” to the consumer must offer at least one service at a reasonable price that allows their subscribers to use any application of their choice without hindrance based on the choice of application.

And, the following language with respect to rule 2 regarding the accessibility of diverse content:

If a sender and receiver make the joint choice to transfer lawful content using an application that is available, that action should not be blocked or hindered based on the content.

If there is not an alignment of interests between sender and receiver, that mismatch will be resolved somewhere in the path from the sender to the receiver. ISPs should be allowed to implement the mediation, but should not unilaterally play the role of policy-makers.

While the proposed rules of the FCC appear to make a clear distinction between applications and services on the one hand (rule 3) and content (rule 1), we believe that there will be some activities that do not fit cleanly into these two categories. We have followed the lead of the FCC in framing our responses in terms of that distinction, but we observe that some further thought about this more basic definitional issue may be helpful in stating clear rules.