

[A Nondiscriminatory Net: The Right Approach](#)

By [Chris Riley](#), January 19, 2010

This is the first in a series of posts by Chris Riley, Free Press Policy Counsel, to summarize the primary policy recommendations made in recent comments submitted to the Federal Communications Commission in its open Internet proceeding. Today's topic: nondiscrimination.

The FCC is in the process of creating a Net Neutrality rule, and front and center in the proceeding is the proposed rule of nondiscrimination – the idea that network gatekeepers should not abuse their control over the pipe by discriminating in favor of or against any online communications. The first of the existing FCC principles defends the right of Internet users to have access to content, applications and services – but it says nothing about the quality of that access.

Under the existing principles, network operators like AT&T and Comcast could speed up their own content, applications and services, and could slow down their competitors' – an especially troubling proposition when the nation's largest fixed Internet access service provider, Comcast, announces a merger with one of the nation's largest content companies, NBC. Network gatekeepers can also use discrimination to impede services like Voice over Internet Protocol and online video – slowing down their bits can render them undesirable to users – and can maintain the stranglehold that the gatekeeper currently has over voice and video services.

But the harms here are not to competition only. Network gatekeepers can also reduce or delay investment if permitted to degrade any service that isn't popular enough, steadily turning the Internet into a walled garden and stifling innovation and evolution in the process. They can also increase their already padded revenue streams – while reducing investment in their networks – if they're allowed to charge exorbitant fees to content, applications and services companies for the privilege of using their terminating access monopoly. Sounds like a great proposition for Wall Street, but it would be the end of the Internet as we know it.

To stave off these dire harms, the FCC has proposed a fifth principle and a rule on nondiscrimination intended to prevent network operators from discriminating against or in favor of different uses of the Internet. The language proposed in the agency's NPRM is clear: "Subject to reasonable network management, a provider of broadband Internet access service must treat lawful content, applications, and services in a nondiscriminatory manner."

Defining Nondiscrimination

Last week, Free Press filed comments with the FCC in support of a strong Net Neutrality rule. Our filing supported the FCC's language – with a big caveat that the terms "nondiscriminatory manner" and "reasonable network management" must be defined correctly. "Reasonable network management" is a tricky issue, and I'll tackle that in a later post. Today's focus is on "nondiscriminatory" and the different options on the table for that concept. The summary version of our position is that a clear and complete rule that prohibits all discriminatory behavior is essential to avoid loopholes; provide clarity for consumers, network gatekeepers, investors and Internet developers; and protect consumer choice, competition and innovation.

Our comments focus on the two most prevalent definitions for a nondiscrimination principle: a clear rule against discrimination subject to reasonable network management, versus a rule that instead prohibits only "unjust and unreasonable discrimination." Sometimes the second option uses other language, such as AT&T's proposed "unreasonable and anticompetitive discrimination" standard, but regardless of the

qualifiers, the proposal is the same: Declare some forms of discrimination to be beneficial, without looking at the network management practice itself.

The first big problem with the “unreasonable discrimination” approach is that it’s utterly unworkable. The goal of powerful network gatekeepers in this proceeding (and, well, every other regulatory proceeding) is to create rules that are made of mush, with giant loopholes that their armies of highly paid K St. lawyers and lobbyists can get through. So their ideal fifth principle is as vague and as devoid of limitations or meaningful standards as possible. That serves their purposes admirably.

What does “unreasonable discrimination” even mean? Is it “reasonable” for a network operator to block BitTorrent or another peer-to-peer program, just because it sometimes uses a lot of bandwidth? That’s in the eye of the beholder. It might seem to network operators that selectively targeting that application makes sense. But for consumers, they’re paying for that bandwidth, and how they use it should be up to them; no network gatekeeper should be allowed to interfere with their lawful use of the Internet, so any discrimination should be unreasonable. And a rule that prohibits only “unreasonable discrimination” is a rule that leaves unfettered discretion to the FCC to choose between consumers and network gatekeepers in each individual case, providing no certainty for any of the affected parties (and no clear standards on which a reviewing court could hang a reversal).

Application Bias

We spend a section of our comments on one particularly insidious idea: application bias. The idea is that a network gatekeeper declares some applications to be “high priority” and others “low priority,” and engineers the network to favor the “high priority” applications. One problem with this over-engineering is that it’s unnecessary. The Internet has performed well without it when network gatekeepers respond to growth in usage by adding capacity. But the biggest problem is that the network gatekeeper is a terrible proxy for the user in determining what is, and is not, high priority. As we showed in a white paper, [*The Hidden Harms of Application Bias*](#), these mistakes can and will harm innovation on the Internet, and will create advantages for incumbents that undermine user choice and competition.

Here’s the bottom line: Any form of discrimination harms some traffic. It slows down some data traveling on the Internet, and speeds up other data. Yes, it’s possible for some forms of discrimination, in some contexts, using some techniques, to create a net positive for the Internet. But there are a lot of qualifiers there, and all of them are important. Any rule that makes a blanket exception to declare some forms of “discrimination” to be acceptable, without any contextual analysis, would harm consumers, competition and innovation. The right approach is not to create a mushy, vague rule with built-in K St. loopholes, but to prohibit all discrimination, and then to evaluate potentially beneficial discrimination through a clear (and standalone) framework of reasonable network management.

[Clear Standards for Reasonable Network Management](#)

By [Chris Riley](#), January 20, 2010

This is the second in a series of posts by Chris Riley, Free Press Policy Counsel, to summarize the primary policy recommendations made in recent comments submitted to the Federal Communications Commission in its open Internet proceeding. Today's topic: reasonable network management.

In my [last post](#), I discussed Free Press' position on nondiscrimination, and why a clear and comprehensive rule without loopholes is essential to protect consumers, competition and innovation. But are there times when discrimination is beneficial and should be allowed? Although the ISPs' network problems are exaggerated, the answer is yes – in the right contexts, when done in the right way. We support the FCC's general idea of allowing for “reasonable network management,” but as we discuss in [our filing](#) for the agency's NPRM on open Internet rules, we remain highly critical of the vagueness of the definition proposed.

The discussion of network controls centers on the issue of congestion, which is often discussed but less often understood. Congestion occurs as a result of very heavy network use, when many users are simultaneously sharing a network resource that has been designed for use by only a few. When congestion occurs, not all of the data passing through a pipe can fit. Congestion sometimes lasts for only a fraction of a second, but sometimes it lasts much longer. Although the Internet was designed to handle congestion without the network faltering, poorly engineered Internet applications or applications that depend on high performance can be temporarily disrupted. The higher the utilization gets within a network, the more frequent and severe the congestion and impact on Internet use.

Previously, network operators dealt with high utilization by increasing network capacity. If too many people used the network at the same time, then it was expanded to accommodate the high demand. This system worked well for the history of the Internet, and it's still working well today.

Delusional data usage

A big talking point of network gatekeepers in the FCC's proceeding on open Internet rules is the idea that Internet use is getting out of hand, leading to unprecedented levels of use and severe congestion. Allegedly, network gatekeepers need to impose similarly unprecedented additional controls to deal with this congestion. This idea is commonly known as the “exaflood,” and it's a delusion. Data usage has been growing steadily for years; network engineers have always been able to accommodate the rapid pace of growth; and there is no evidence to support any change to this pattern. Talk about a solution in search of a problem!

That said: Even a properly engineered network will experience sporadic, mild periods of high utilization and congestion. So appropriate network controls to deal with periods of congestion can be reasonable. Similarly, network controls that deal with spam or viruses or [denial of service](#) attacks can certainly be reasonable. And the FCC should ensure that network operators' ability to impose “reasonable” network controls can coexist with consumer protections in this proceeding.

Clear guidance from the FCC

So, where does that leave us? The FCC needs to offer clear guidance as to what “reasonable” means, so that network operators can better understand what actions might get them in trouble, and consumers can be assured that this “reasonable” framework won't just rubber stamp anything the network operator

wants. But rather than be clear, the FCC's suggestion was circular: "Reasonable network management consists of practices which are reasonable."

Admittedly, this isn't an easy question. The FCC has two conflicting goals: Be clear, but allow for good behavior. But the right solution isn't that hard. It's similar to frameworks that have been adopted either voluntarily or by regulation in Japan and Canada. To be considered reasonable, network controls must have a good reason for their existence, and they must not harm anybody unnecessarily. Put in other terms, this is a two-part test for "purpose" and "means."

Public interest purpose

The first part of the test is: What is the purpose of the practice, and why should we consider it valuable? In our comments, we propose that the purpose be a "public interest purpose" – something which on balance serves the public interest, not merely short-term parochial interests. The purpose should also be real, not hypothetical. For example, congestion management can be a public interest purpose, but only if the network operator can demonstrate that congestion is occurring or at least likely because of high utilization. Going back to my [first post](#): Any discrimination is inherently harmful to some Internet traffic. Thus, if the purpose isn't real, the discrimination at issue is unwarranted and unreasonable. Practices should not be approved absent data showing that the intended purpose is not merely hypothetical.

Time, geography and proportion

If the purpose is real and valuable, the next question is, what means are used to achieve that purpose? We break this question down into three parts: time, geography and proportion, and I will use congestion management as my hypothetical purpose. "Geography" says that if you demonstrate high utilization in a service area somewhere in downtown San Francisco, you should not be using a congestion management practice in Boston. "Time" says that if you demonstrate that your network experiences high utilization between 7 and 9 p.m. in a local area, you should not be using a congestion management practice at noon. And "proportion" says that you can't block all uses of one application even if some of those uses are contributing to congestion. You can't discriminate just because it's the easy answer if you have a more appropriate remedy available.

With a good two-part purpose and means test in place, the FCC can evaluate network controls that violate the open Internet rules on a case-by-case process, to separate the bad actors from the good. Clear standards for what counts as "reasonable" will limit the FCC's arbitrary discretion, and give the reviewing courts (and the public and Congress) something to evaluate. But a definition that says "reasonable is that which is reasonable" is nothing but Swiss cheese.

[FCC: Apply the Same Rules to Wired and Wireless Networks](#)

By [Chris Riley](#), January 25, 2010

This is the third in a series of posts by Chris Riley, Free Press Policy Counsel, to summarize the primary policy recommendations made in recent comments submitted to the Federal Communications Commission in its open Internet proceeding. Today's topic: wireless networks.

I began this series discussing [nondiscrimination](#) and [reasonable network management](#), which lie at the heart of the FCC's open Internet proceeding. In both of these posts, I referred generally to "networks," and I didn't mention cable, DSL, FiOS or wireless. This was intentional – there should be no distinctions in open Internet policies based on the access technology, and in our FCC filing, we argue that no Internet access service provider should be treated differently from any other.

With all Internet access services, a consistent, clear, reasonable network management framework will permit any legitimate practices by network operators to deal with congestion and other network problems that may arise.

In our comments, we call for the FCC to recognize that there is only one Internet, regardless of the technology used in the last mile, be it cable, DSL, fiber or wireless. Each Internet access service provider is a gateway to the same Internet, and commerce and speech flow from one technology to another. Consumers don't distinguish between access provided via Internet devices, applications and network technologies. The same pro-consumer rules must apply, regardless of the technology used in the last mile of the connection.

Long gone are the days when Internet users were chained to a fixed desktop computer that could only be used with one network connection. Laptops are widespread and rapidly replacing desktop computers for a generation of technology users. Many of these laptops can readily switch from an Ethernet cable to a WiFi network to a 3G wireless connection. Many popular phones have both WiFi and cellular radios, and can run the same applications and access the same content as laptops. And netbooks, tablet PCs and the newest generation of smartphones like the iPhone and the Nexus One blur the lines further. From the Internet user's vantage point, there are no longer categories of fixed and mobile devices, just devices.

Similarly, in the not-too-distant future, 4G wireless networks may well be an effective substitute for many forms of wired connections, such as DSL lines. As with devices, eventually the lines between mobile and fixed Internet access services will blur. (In fact, wireless carriers often argue that even current generation 3G networks are substitutes for fixed connections.) Internet users eventually may not distinguish between wired and wireless networks, either in devices or their expectations for performance and access to content and applications.

Future-proof rules

Open Internet rules adopted in this proceeding must be future-proof, and thus must not distinguish between wireless and wired networks. Failure by the FCC to apply open Internet rules to wireless networks would permit great harm to wireless consumers, and skew innovation and competition in the broadband services market for years to come.

Admittedly, many modern wireless networks do not work as well as wireline networks, though many of the limitations are the result of bad business decisions and a poor market, rather than something specific to wireless [<http://www.freepress.net/node/76177>]. But the solution to these problems isn't to allow wireless companies to engage in anti-consumer and anticompetitive behavior. The solution is to

encourage investment and competition through good public policy. It's to encourage wireless providers to build more towers with better backhaul connections and to lower prices for consumers. And rather than being harmful to this goal, open Internet rules actually promote competition – they will encourage investment and consumer-driven behavior in the wireless market, just as they do in the wired market.

Also, when facing temporary problems with congestion and other resource limitations that cannot be addressed through investment, wireless network operators have the same opportunities as wired network operators – the freedom to engage in reasonable network management to address public interest problems using proportional means. In my proposal from [my last post](#), I said that network controls should be proportional in the context of the network and the stated purpose. A “proportional” test can accommodate any limitations, including those that arise from spectrum or other problems inherent to the wireless industry, as well as problems that are shared by all Internet access service providers. DSL is not like cable, and cable is not like fiber, and these technologies also have unique characteristics and limitations – wireless is not magically different. Whether for wireless or cable or DSL or fiber, case-by-case processes for evaluation allow plenty of room for the network operator to demonstrate limitations that justify additional network controls.

Freeing wireless devices

Our comments also address another major issue for wireless raised by the FCC: devices. In 1968, the FCC broke open AT&T's closed phone network in its landmark Carterfone proceeding. The result was a flood of innovation in phones and phone add-ons, and a wealth of economic and consumer benefits. Here, the FCC has essentially proposed extending the ideas of Carterfone to wireless: Let consumers connect any phone or any device that works to the network, to use the service they are paying for with the device of their choice. We strongly support the right of Internet users to attach any device of their choosing to their wireless network. You can use a Dell or Apple laptop with your cable modem; why should your wireless carrier be allowed to prevent you from using the Samsung, Motorola or Blackberry of your choice?

Sure, there are technical limitations that must be dealt with. Some cell phones use GSM radio technology; others use CDMA, and only a few can use both. But if a device is compatible, and if the FCC certifies the device as non-harmful, the service provider should not stand in the way, and shouldn't be allowed to impose more limitations. Wireless carriers should not be allowed to refuse device connections, or to impose certification barriers on devices that the FCC has evaluated and declared safe for use. And connection and service use should be allowed whether a user wants to connect through a phone, a laptop or a laptop through a phone using a tethering application.

On all the wireless issues in this proceeding, including nondiscrimination and devices, the role of the wireless service provider should be to provide and maintain the wireless service, not to micromanage the use of that service and artificially control the choice and freedom of users.

[Clear and Complete Disclosure, Without Loopholes, Is the Only ‘Reasonable’ Way](#)

By [Chris Riley](#), January 26, 2010

This is the fourth in a series of posts by Chris Riley, Free Press Policy Counsel, to summarize the primary policy recommendations made in recent comments submitted to the Federal Communications Commission in its open Internet proceeding. Today’s topic: disclosure.

Consumers have a right to know what is going on with their Internet access service. The FCC under both Democratic and Republican leadership has long held disclosure to be among the most important values to an open Internet. The agency has proposed making network operators disclose their interference with subscriber Internet use.

As Free Press points out in [our filing](#) with the FCC’s NPRM on open Internet rules, we support mandatory disclosure wholeheartedly, although it cannot take the place of nondiscrimination and other open Internet rules. Also, the specific language proposed by the FCC has a number of loopholes that, in practice, could render disclosure meaningless and leave consumers subject to hidden and harmful network controls.

We’ve already seen the harms of secret Internet interference. For years, Comcast blocked the use of peer-to-peer services. Comcast used a [DPI-based system](#) that made its blocking look like an ordinary network action, so that subscribers couldn’t tell what was happening. It took months of expert analysis, and even then, the public (and the commission) could only determine the broad details and not specific triggers (which remained unknown until Comcast made its [FCC-mandated disclosure](#)).

The current state of voluntary disclosure of network interference is simply nonexistent. Comcast disclosed its [“new system”](#) when the FCC ordered it to, but the company has said that it can change its system any time it wants, without telling anyone. And no other provider has even given that much information.

Full disclosure

Internet subscribers pay a substantial amount of money for service every month, and in return, they are presented with vague and overbroad “terms of service” in which network operators say that they will block any Internet communications they disagree with, and disconnect users for behaving in any manner they find objectionable. In the FCC’s “truth in labeling” proceeding, we’ve [asked the agency](#) to fix this – to require network operators to provide consumers with meaningful and accurate information about the real restrictions and limitations on their service. But in addition to these changes, network operators should meet a higher standard of disclosure when they interfere with subscribers’ use of an open Internet access service.

As in the Comcast case, harm that is not obvious is still harm. Internet users may experience problems without knowing they are caused by network operator interference – especially when some other Internet uses work fine. With the ISPs quietly interfering behind the scenes, users will blame the application or Web site they’re using, rather than the real culprit, and may choose to stop visiting that site or using that application. Innovation will pay the ultimate price.

In response, the FCC has proposed a weak and ineffective standard. The FCC limited the required disclosures only to information that is “reasonably required” by Internet users “to enjoy the protections” of the open Internet rules. Furthermore, the FCC’s proposed disclosure standard is “subject to reasonable network management.” This standard is too vague to protect Internet users and to secure the open Internet.

Closing loopholes

The first loophole is the idea that disclosure could be “subject to reasonable network management” and that its purpose is “to enjoy the protections” of the open Internet rules. Both of these phrases imply that any network management practice that is “reasonable” need not be disclosed. If a network management practice is reasonable, then Internet users would be subject to it, so its disclosure would not come under the “protections” of the rules and may not be required under a “reasonable” exception. This would be an absurd result. It would require disclosure only where the disclosed activity is, in fact, illegal. The FCC must close this loophole by eliminating both of these qualifiers.

The second loophole, although not as severe, is the idea that ISPs need to disclose only information “reasonably required” by Internet users. This gives far too much discretion to network operators, which have incentives to restrict the awareness of users and competing content and application developers. Empowering network operators, which alone have information about network interference, to determine what of that information can and cannot be disclosed does not fix the asymmetries. The FCC seems concerned that disclosure obligations should be “minimally invasive” – but any network company with decent engineers will have detailed information at its fingertips on the network interference at issue, and likely will have plain language characterizations of the practices already written, for use in other business divisions such as customer support and management. It’s not invasive to require a company to turn over information it already possesses, or to make minor revisions to existing information to make it suitable for public dissemination.

The right solution

The FCC’s approach to disclosure seems misguided. It searches for a magic level of detail in disclosure that will satisfy all needs. Such an approach is utterly unworkable – any disclosure that a substantial population of Internet users will readily understand will fail to provide the details needed by a competing content or application provider (and by high-level users) to make the most effective use of the Internet access service. The right solution, as we propose, is a two-level disclosure.

The FCC should require a high-level disclosure of basic information about actual network interference that is prominent on Web sites and clearly available at the time of service inquiry and on bills.. The FCC should also require a more detailed level of disclosure to any interested third party – not merely a content or applications provider – containing detailed information on trigger thresholds and specific methods of interference. This information should be provided at the time an Internet user signs up for service, and should always be available and up-to-date on the Internet access provider’s Web site.

As part of its open Internet rules, the FCC should require clear, complete and detailed disclosure of all network interference, without any loopholes. Internet users deserve nothing less.

[The Gatekeepers' Discrimination Delusion](#)

By [Chris Riley](#), January 28, 2010

This is the fifth in a series of posts by Chris Riley, Free Press Policy Counsel, to summarize the primary policy recommendations made in recent comments submitted to the Federal Communications Commission in its open Internet proceeding. Today's topic: business models.

Investment is one of the most heavily debated issues in the open Internet proceeding. What regulatory paradigm will lead to the most private sector investment? Although the government plays a major role in building our national communications infrastructure, private investment supplies a lot of the capital to lay the wires and build the towers, investment that often pays off handsomely for investors. There's a lot of rhetoric bandied about here, including the classic trope that "all regulation discourages investment," which does not accurately reflect how communications markets work.

When thinking about investment in the Internet, it helps to separate out two pieces – investment by network operators in the underlying facilities, and investment by content and applications companies in the Web sites and software tools that we use. The Internet relies on both of these investments – without either, we are worse off. There's also a bit of a zero-sum game at work here, because discrimination by network operators, even if it could increase investment in facilities, [would reduce investment by the content and applications companies](#). But the main argument made in our filing with the FCC is even more direct: Network operators won't have an incentive to invest in their facilities if they are allowed to discriminate, so open Internet rules will not discourage investment. The "discouraging investment" argument is nothing but a scare tactic to dissuade the FCC from passing strong rules to protect the open Internet.

The fallacious argument offered by network operators is this: Discriminatory practices permit them to invest in "new business models" that will increase their revenue. They seem to suggest (despite the absence of a competitive market) that rather than pocket this extra revenue or distribute it in the form of shareholder dividends, they will instead use it to increase their investments in network infrastructure, at least by enough to offset any reduction in investment by content and applications companies as a result of discriminatory practices.

Discrimination Delusion

A close look at the market and the business models reveals this line of reasoning to be nothing more than a great "discrimination delusion." The true purpose of engaging in discriminatory practices is not to create new revenue streams, but to protect existing ones – the traditional phone and cable services that now face competition from innovative Internet-based voice and video offerings.

There are three possible types of business models based on discrimination, and none of them would make the Internet a better place. Even where these models can successfully raise network operator revenue, the increase is unlikely to be substantial. These plans also produce powerful incentives to reduce investment in the network, and little to no incentive to expand network coverage or lower service price.

Pay for Play

Under the first category of business models – "pay for play" – network operators charge third-party content and applications providers fees above and beyond normal transit costs to permit their traffic to be routed through their network. Recalling an old and well-known quote from former AT&T CEO Ed

Whitacre, this business model will prove ineffective at raising revenue because so much of the Internet's value to consumers lies in its content and applications, not in its underlying facilities.

Insofar as its value lies in the content that travels through it, the Internet resembles cable television systems, where the network operator actually pays the content company for the right to carry the content. The cable model of networks paying for content, rather than the "pay for play" model of content companies paying for carriage, is far more sensible, realistic and likely to prevail. In fact, the pay-for-content model is already in use on the Internet today – ESPN offers Internet content, ESPN360, exclusively where the network operator pays a premium for carriage. So in the long run (if not the short), a "pay for play" model will not result in a new revenue stream for network operators, and may in fact cost them money and thus raise users' Internet connection costs even more.

Pay for Priority

The second category, "pay for priority," involves selling priority treatment for some content or applications. As we have argued [elsewhere](#), prioritization only works in an environment of widespread congestion, where speeding up some traffic slows down other traffic. A priority service has value only in this environment, but additional investment by network operators would reduce congestion, and thus the potential value of the add-on service. So, network operators able to offer "pay for priority" services would have an incentive to avoid investment, to maximize the potential value of such services.

Additionally, the network operator can only offer a few prioritization deals. The more deals that are made, the less the priority given, because more content would move through the same narrow lane, slowing everything down. An increase in the number of priority customers thus lowers the value of a potential prioritization service, and lowers the revenue that can be made from each deal. After a limited number of these deals, no more can be made without rendering them all worthless – so there's a tight cap on the amount of revenue this strategy can generate.

In other words, the financial benefits of "pay for priority" models are dubious and will discourage rather than encourage investment in network facilities. It just doesn't add up to a better network or to more investment.

Vertical Prioritization

In the final category of business models, "vertical prioritization," a network operator prioritizes its own content and application offerings, avoiding any deals with third parties. Under this model, the network operator has no new revenue streams. Instead of new revenue, vertical prioritization allows the network operator to insulate its existing offerings from competition.

Insulating old business models from competition is a recipe for stagnation and higher consumer prices, not for investment in network facilities and economic growth. Network operators can even reduce their investments and allow their network to grow more congested, while maintaining or raising rates for their affiliated content and applications (and especially for their legacy phone and cable businesses), because then theirs are the only offerings that work.

Ultimately, abandoning the open Internet would create an environment where network operators can more easily and will likely reduce network investment. Closing the Internet would drive network operators to manage and seek to wring profit from scarcity, not to invest in capacity to meet Internet user demand. We would face a less robust, and less interesting, Internet – and a smaller financial pie of which the network operators would have a much larger share.

[The Open Internet Is in the Public's Interest](#)

By [Chris Riley](#), January 29, 2010

This is the sixth and final post in a series of posts by Chris Riley, Free Press Policy Counsel, to summarize the primary policy recommendations made in recent comments submitted to the Federal Communications Commission in its open Internet proceeding. Today's topic: why open Internet rules are in the public interest.

In addition to our assertion that protecting the open Internet will not harm network investment, the subject of my [last post](#), our [filing with the FCC](#) makes many other compelling and detailed arguments in support of strong rules to protect Internet openness. These arguments center around one idea too often lost in this debate: The Internet is not AT&T's Internet, nor Google's Internet, nor Free Press' Internet – it's the public's Internet. The Internet is not a widget, but public infrastructure upon which ever larger parts of our society are being built. There is no economic substitute. The Internet is a public resource, and a public good.

Even if Google and Verizon can find ways to get along, the aim of communications policy should not be to referee industry disputes, but to identify and establish policies that best promote the public interest. And to defend the public's interest in the public's Internet, it must be kept open. The FCC must adopt strong and clear rules, without loopholes, that establish lasting protections for the current environment and for the future.

- Keeping the Internet open is in the public interest because it will promote competition among network operators, among content producers, among application developers and among service operators. If network operators are unable to maintain their market dominance through anticompetitive protections of their existing voice and video business models, they will have more incentive to offer a quality, inexpensive broadband Internet access service to gain and retain customers. We need more than open Internet rules to turn our nation's broadband duopoly into a competitive market, of course. But preserving the open Internet is a necessary, if not sufficient, condition.
- Keeping the Internet open will increase deployment and spur adoption. It will promote investment and competition for networks and for applications and content. Investment drives more rapid and higher capacity broadband deployment into unserved and underserved communities. Competition lowers prices and improves service quality for all consumers. Investment, competition and an open platform for innovation will increase the diversity and availability of online content and applications, increasing the Internet's value for would-be subscribers and encouraging adoption. As some have pointed out, the alternative is essentially the telecom equivalent of [trickle-down economics](#).
- Keeping the Internet open will create jobs. From 1996 to 2004, when the telephone companies operated in an environment of regulation and competition, their revenue and jobs curves tracked: When they made more money, they hired more people. But with deregulation (starting in 2005), these curves diverged. Since then, revenue at AT&T, Verizon and Qwest has steadily increased – yet the companies have shrunk, firing more employees and reducing relative investment levels. It's time to reverse this trend. Open Internet rules will promote competition and investment, and therefore will encourage network operators to hire more people and build bigger networks, rather than enable them to continue to increase their profits through discriminatory behavior while reducing network investment.

- Keeping the Internet open will promote consumer choice. Although rules to preserve the open Internet will not give most consumers more than one or two choices for true broadband Internet access services, these rules will promote competition and investment that will help us eventually reach the goal of having many choices available. And in the meantime, they will protect one area of robust consumer choice on the Internet – the online market for content and applications. Failure to pass strong open Internet rules, or permitting loopholes in those rules, would permit network gatekeepers to leverage power into this space, and to stifle competition and consumer choice in the one Internet sector where it currently exists.
- Keeping the Internet open will promote innovation. Consumer choice and low barriers to entry for content and applications online create a market where no incumbents can rest on their laurels. We have seen the rise and fall of many major Internet companies that were too slow to innovate and adapt to the changing needs of their users. Protecting the open Internet will preserve this dynamic – but closing it would allow incumbents to use their tools and legal muscle to protect their business models against new and more innovative competitors. Comcast should not be permitted to block online video providers that compete with its cable television service, just as AT&T should not be allowed to interfere with VoIP software that threatens its voice business.
- Keeping the Internet open will promote free speech. Preserving the open Internet will ensure that users have few barriers to public expression, communication and association, through ever newer and more innovative media like YouTube, Facebook and Twitter. No gatekeeper should be allowed to put a tax on the speech of Internet users.
- Keeping the Internet open will promote democracy. The open Internet is not just a forum for speech and commerce; it is part of democracy's essential infrastructure. Governments connect with citizens through the open Internet, and the resulting two-way communications and public participation are reinvigorating and reshaping our democracy in new and exciting ways. We must not approach communications policy as though the Internet were just another economic product or service, or we will undermine that future.

If that sounds like a lot of reasons, well, it is. And the public has spoken – as of now, more than 150,000 people and dozens of organizations have told the FCC that they support strong rules to protect the open Internet. [Join us](#) as we fight to keep the Internet open, and to shape communications policy in the public interest.

[Market Failure in the Wireless Industry](#)

January 21, 2010

Posted by Chris Riley

There are two major signs of market failure in the wireless industry right now: Service quality is often poor, and service price is always excessive. Brand a scarlet letter on AT&T and Verizon, the two biggest incumbents.

Too often, these problems are blamed on others –sometimes even on the public -- but the fault lies with the network operators alone. AT&T's wireless network just isn't very good, and it's not because of insufficient spectrum or too many iPhone users. It's because of market failure, the result of an environment where AT&T doesn't have to invest aggressively to gain and keep customers. At the same time, wireless carriers routinely raise the prices and fees associated with their services, because they can safely do so without losing business, especially when the pricing practices of the two industry behemoths, AT&T and Verizon, are exactly the same for voice, data and text services.

This pattern of poor service quality and high service price, despite above average operating margins in poor economic climates, should be a clarion call to the FCC, which is currently preparing its fourteenth report on competition in the wireless market. All is not well, and something needs to be done.

Failures to build good wireless networks

Washington Post reporter Cecilia Kang recently [described](#) service quality problems with AT&T. At the Consumer Electronics Show in Las Vegas, Kang interviewed several iPhone users reporting major problems with connectivity – from 15-minute delays in checking e-mail to complete nonfunctioning of applications. At the bottom of the article is a gem -- a few quotes from Dick Lynch, chief technology officer for Verizon Wireless. Lynch reported no problems with Verizon, saying, “What we have done and continue to do is have a buffer of capacity above what our demand is at any given point in time. When we see consumer demand begin to feed into that capacity, we scurry out there to add more capacity.”

We've filed [comments](#) (p. 17-25) in the past demonstrating AT&T's failures to build its network to accommodate iPhone users. We pointed to evidence that AT&T anticipated the growth in usage – some sources say they expected it to be even greater – and ultimately chose not to invest enough to meet that demand. This isn't a problem with the iPhone. And it isn't a spectrum crisis. It's an “AT&T's network sucks” crisis. And consumer studies are backing this up, whether from [Consumer Reports](#) or from [Zagat](#).

Now more people are realizing what consumers and consumer interest groups have long been saying. [PC World](#) details an analysis of AT&T's network and the company's business behavior, courtesy of TownHall Investment Research. According to TownHall, AT&T gets 57 percent of its revenue from wireless, but only allocates 34 percent of its capital expenditures for maintenance and buildout of its wireless network. TownHall's conclusion: AT&T needs to invest a lot more money in its wireless network if its customers are going to have a satisfactory experience.

As we move toward next-generation wireless networks over the course of the next decade, will we eventually need more spectrum? Probably, yes – and it's wise for the FCC and NTIA and Congress to look now for ways to get spectrum down the road when it's needed. But the problems of service quality right now aren't a result of spectrum limitations but of bad business decisions.

Failures to fairly charge consumers

The market doesn't look any better on the price side. Yes, AT&T and Verizon recently lowered their prices for unlimited voice service by \$30 per month, but after steady talk of raising rates on data usage and moves by both carriers to [impose data plans](#) on more users, it's clear to reporters and policy folk alike that the adjustment isn't so much a price cut as a [price reallocation](#).

The price shift for high-end services affects only a few users, and the ones who already pay the most for voice service. The rest of us, who choose not to spend hundreds of dollars per month on our service and who don't use that many voice minutes, won't see any benefit. In fact, our bills will [go up](#) when Verizon and AT&T raise data rates. Total revenue for the companies is projected to [increase](#), not decrease, as a result of these adjustments. But total revenue might increase for Verizon and AT&T even before they raise their data rates, because consumers like unlimited service plans, and some who are currently purchasing the next level down of service will upgrade even unnecessarily, to avoid the risk of overages.

In fact, as we have argued before, wireless network operators set the individual use price of voice minutes, data transfers, and text messages artificially high, to deliberately inject fear of overages into the minds of consumers, and to get them to pay for more than they need to. As with other pricing tactics, this is not a charitable move by the incumbent carriers. It's just a ploy, plain and simple.

And, if Verizon and AT&T can afford to chop \$30 off their high-end service price, why can't they reduce it even more? The [analysis of service plans](#) from BillShrink shows exactly identical service offerings and prices from Verizon and AT&T. Game theory tells us that as long as Verizon and AT&T don't underbid each other, both are better off, and as soon as one breaks from the trend and underbids the other, the resulting price war will hurt their 30 percent profit margins – though it certainly would help consumers. The FCC can make this market better, but it has to look at more than spectrum to do so.