

Valuing Network Neutrality

“Net Neutrality” has been the subject of heated debate over the last two years. A business-school professor and a self-proclaimed “network geek” find that neutrality promotes competition and innovation¹

By Mark Gaynor ■ *Boston University* and Scott Bradner ■ *Harvard University*

Should telecommunications carriers provide neutral networks – “dumb pipes” that treat all bits and bytes equally – or should they bias the network in order to privilege certain services or content? Economic models have shown that neutral networks create more revenue under certain circumstances. This article expands on those models to show that the network’s value depends both on how neutral it is and how uncertain the market for services might be.

Today’s Internet technology is changing so fast that network providers can’t predict which services individual users will want or what value the users will place on services and content. That means users can’t select a single service or content provider that is the best match for every user. When network providers bias the network toward specific services or content, they at the same time give up on the opportunity to offer some services or content that some users would have paid extra for. Our analysis shows that a non-neutral network ultimately stifles the creation of a rich competitive ecosystem of infrastructure providers, service providers and content providers.

What Is Net Neutrality, Anyway?

While the term “net neutrality” is new, the concept is not. An 1860 federal law subsidizing the construction of a tele-

graph line from the West Coast to the East Coast contains a sentence reading, “Messages received from any individual, company, or corporation, or from any telegraph lines connecting with this line at either of its termini, shall be impartially transmitted in the order of their reception, excepting that the dispatches of the government shall have priority.”

Columbia Law professor Tim Wu, who is credited with coining the term “net neutrality,” calls it “preserving Darwinian competition for every conceivable user of the Internet so that only the best will survive.”

The creator of the Web, Tim Bern-

ers-Lee, defines network neutrality as follows: “If I pay to connect to the Net with a certain quality of service, and you pay to connect with that or greater quality of service, then we can communicate at that level.”

While these definitions vary, the underlying theme is fairness and competition among providers of services over the Internet.

There are “hard” and “soft” definitions of network neutrality. The “hard-line” definition for the Internet is that the network should have only one class of traffic:

best effort. A “softer” approach would allow different tiers of traffic handling as long as quality of service (QoS) is sold in a fair and nondiscriminatory fashion to all users and service/content providers. We argue for soft network neutrality for several reasons. First, it seems unreasonable to micromanage the services available on privately built networks. Second, the soft definition of neutrality accommodates QoS products from vendors such as Cisco and Alcatel/Lucent, as well as protocols supporting QoS, such as MPLS and DiffServ, that have been developed by the Internet Engineering Task Force.

Yet another reason to allow a “soft” definition of network neutrality is that best-effort transport might not be good enough as content and services become more dynamic, and as VoIP and video traffic grows. However,

there is still debate about this: Industry pundits such as David Isenberg argue that a better approach would be to build networks where the worst QoS works well enough for real-time traffic (in other words, to substitute additional bandwidth for service management intelligence). This issue is discussed more rigorously by Yuksel in a model describing how much additional bandwidth is required to meet a given level of QoS.²

Under any definition, neutral networks do not bias traffic based on content, source, destination, or any other

Today’s Internet technology is changing so fast that network providers can’t predict which services individual users will want or what value the users will place on services and content.

attribute of the data or metadata about the data in ways that are undesirable to any end user, content provider or service provider. This is fundamental to providing network users with unbiased choices in an environment conducive to innovation and investment.

Network neutrality boils down to providing users with an unbiased choice of services and content, which our model implies will promote greater user satisfaction as well as opportunities for innovators and entrepreneurs.

But neutrality is an unrealized ideal. Examples of non-neutrality abound in the Internet, cable TV networks, and the Public Switched Telephone Network (PSTN). In the cable TV network, a good example of non-neutrality is Comcast's Video on Demand service, where users' only choice of content is from Comcast.

Network Neutrality and Architecture

Network neutrality is not always related to the architecture of the network. Networks with similar infrastructures can be managed neutrally or non-neutrally.

Centralized control (such as the PSTN has) makes it easy to control access to content and services. However, even centrally controlled networks can be neutral if the governance of the network allows it. One example of this is the infrastructure that is being developed in France for wireless application protocol (WAP) Internet services via GSM or GPRS 3G mobile phones. The French government has published conditions promoting network neutrality for WAP mobile wireless Internet services. French courts have ruled that users must have a choice of WAP service providers, and that mobile devices must allow users to easily change the default WAP gateway to a gateway of their choice.

While it is still too early to evaluate the outcome of these French regulations, we believe they will foster a competitive environment in which users have a choice of gateways and portals in the emerging market for mobile Internet services.

For networks with a more distributed architecture, such as the Internet, which is made up of a collection of auto-

When network providers bias the network toward specific services or content, they at the same time give up on the opportunity to offer some services or content that some users would have paid extra for.

nomously owned and managed networks, it is more difficult – though usually not impossible – to bias traffic end-to-end.

True end-to-end infrastructure, such as the early Internet, is neutral simply as a by-product of its architecture. The early Internet, with its core best-effort service model, was completely application-unaware. This architecture enforced neutrality because the simple network infrastructure was ignorant by design of the content or services it was transporting.

But today's Internet is not an end-to-end network. Many organizations use traffic shaping based on information above the IP layer, routers that inspect packet information above Layer 3, and packet classification. These functions are all departures from the original Internet architecture. While it would still be possible to manage the Internet as a neutral network, today's ISPs can and do filter, alter, and block traffic to and from their customers, as well as bias traffic flowing to and from other ISPs.

Some types of traffic filtering are desirable – for example, it is important to block malware traffic and Distributed Denial of Service (DDOS) attacks – but other types of interference, such as slowing or blocking real-time VoIP traffic (see the FCC case on Madison River Communications, Order DA 05-543), are not.

Examples of non-neutral Internet service include ISPs blocking traffic based on port addresses, or providing higher-quality service to application providers they have partnered with (say, Google) than to competing services, or biasing their own VoIP service over competitive VoIP service providers such as Vonage.

What Is Market Uncertainty?

Earlier, we said that the value of net neutrality depended on the degree of

uncertainty in the market for Internet-based services. So, what is this market uncertainty? In brief, it is the inability of service and content providers to predict what users will like and how users value the features of a service or the selection of content. When market uncertainty is low, providers can predict the value of their offerings; high uncertainty means that providers must guess.

One reason that market uncertainty exists is that users often don't know what they want until they see it. Consider the format standards battle between VHS and Sony's Betamax videotapes; the proliferation of hundreds of features for voice calling in the PBX market; the surprising success of the "Hello Kitty" service in NTT's i-mode wireless network; and the innovation of the Web itself.

These examples all illustrate situations where vendors and service providers failed to understand their markets. Sony bet big on its Betamax standard, which it believed to be technically superior, but the market decided otherwise. In the early 1980s, PBX vendors experimented with many voice features in the newly developed software-controlled PBXs, and users selected features they wanted most, such as caller ID and voicemail.

Nobody predicted in the early 90s what the Web would become today, or the impact it would have on society – yet in 10 years the Web emerged as a requirement for modern commerce. When Netscape started developing the first breed of Web browsers, there was extreme uncertainty. Users had no idea what they would do with browsers and, thus, what features they would like, and vendors had no idea what services would become popular. The Internet today is far different from the predictions of early 1990s, which illustrates the high level of

While the term “net neutrality” is new, the concept is not. An 1860 federal law subsidizing the construction of a telegraph line from the West Coast to the East Coast contains a sentence reading, “Messages received from any individual, company, or corporation, or from any telegraph lines connecting with this line at either of its termini, shall be impartially transmitted in the order of their reception, excepting that the dispatches of the government shall have priority.”

market uncertainty that exists in network-based services, and the way users’ preferences evolve with the technology.

Knowing whether market uncertainty is high, medium or low is important to our model. There are several ways to estimate market uncertainty, including the ability to forecast the market; the emergence of a dominant design; and feature convergence. Network services with low market uncertainty – such as traditional voice services – are differentiated by price, not features: all service providers understand what users want and providers all offer essentially the same set of features. However, as market uncertainty grows, user selection criteria migrate from cost to the feature set or available content that best matches the user’s wants.

With low market uncertainty and price-based competition, it is unlikely for any service provider to win big because of low margins. But high market uncertainty and feature-based competition allow service providers to charge more for more successful ideas, allowing the possibility of capturing a large potential market and winning big. For example, Apple guessed right with the features of the iPod and has been able to charge a premium over what other providers charge for their competing products while still dominating the market.

The Relationship Between Market Uncertainty and Network Neutrality

When market uncertainty is high, a non-neutral network that provides users with

only a single choice of services selected by the network operator will, at best, realize an average revenue for services, content, and transport. For a similar network managed in a neutral fashion – that is, where users have many choices for services and content – revenue is expected to be above average because the benefits to users will be higher and the user will be more willing to pay more for the network connectivity. The reason is that a single set of services from a single service provider is unlikely to match users’ wants as well as a group of independent service providers would. A single transport service provider will never develop or select as many innovative services as a group of heterogeneous and independent providers are able to do.

On the other hand, when market uncertainty is nonexistent or low, there is little to be gained from a neutral network. The commodity nature of the services implies that most users are happy with the one choice offered by the network provider and would be unwilling to pay more for the network. Thus, market uncertainty is a critical factor in determining the value of network neutrality.

We can demonstrate this by looking at two extreme situations:

- A completely non-neutral network where users only have a single choice of bundled services at a fixed price.
- A very neutral network where users have many choices for all service and content providers available at many different prices.

Even though these two extreme cases are not realistic, they illustrate the limitations of a non-neutral network versus the potential of network neutrality.

But first, let’s look at where the value of a network comes from.

How Investors Value the Network

Let’s consider a transport network such as broadband cable or DSL. The value of the network to the investors who are funding it is based on its profits. The profits equal the network’s total revenue, minus its total costs. That is, the total value of the network is

$$\begin{aligned} & \text{transport revenue} \\ & + \text{service revenue} \\ & - \text{cost of transport} \\ & - \text{cost of services} \end{aligned}$$

We assume that the costs will be the same whether the network is neutral or non-neutral. (In fact, non-neutral networks cost somewhat more because they are more complex to manage.)

How Users Value the Network

For users, the network’s value is more subjective. But we can assume that the value for each user is at least equal to the price for connectivity plus the prices for any other network-based services and content, or else the user wouldn’t subscribe.

We can also assume that the connectivity itself has no intrinsic interest – no one would subscribe to a network that didn’t offer access to any network-based services. Users want the services and content that they can access over

the network. Of course, they are willing to pay something for connectivity, but that something depends on how highly they value the services and content. So the total value of the network itself to users is proportional to the value they place on the services.

These services won't be equally valuable to each user – for example, I might enjoy surfing the Web, while you might prefer watching video. But when there are a large number of users, the values they place on each individual service and on the service bundle as a whole will follow the familiar “bell curve” shown in Figure 1. (In technical terms, each user is willing to pay a random amount that is normally distributed for any randomly selected group of services that compose the bundle.)

In Figure 1, the horizontal axis represents how much a user would be willing to pay for a service. The central point represents the price that 50 percent of potential users would be willing to pay. Around the middle, within one stan-

dard deviation from the mean, are the 68 percent of users that are moderately happy with the service. The users on the far right are very happy with the service and value it far more than the average user does.

When market uncertainty is low (well-established services, “commodity” pricing), the users agree about the value of the network, the provider understands what the users think, and the bell curve is narrow. But when uncertainty is high (new services, nonstandardized features), neither the provider nor the users have a good idea what the service bundle will be worth to users, so the bell curve is likely to be spread out.

Case 1: The Walled Garden

The extreme non-neutral network case is called a “walled garden” business model: a single provider for network connectivity, network services, and access to content. Since the provider has limited resources, it must choose a finite set of services to include in its standard

bundle – say, a single solution for each service such as e-mail, VoIP, and video. The provider then sets a price that covers both services and transport, and keeps 100 percent of the revenue for itself. Users have only one choice: Either pay the price the service provider is asking for the bundle of services it selected, or not.

No Market Uncertainty

When there is no market uncertainty, services are commodities sold strictly by price. All service providers know exactly what users want, and will have a reasonable idea what they will pay for it. One hundred percent of potential users will subscribe to the services offered by the single service provider if the provider guesses correctly about the price. This is not far off from the voice service market. Until recently, virtually all users were willing to pay the regulated price for wired phone service via the PSTN. Similarly, voice-only cellular services have captured a very large market percentage at commodity pricing structures that have little variance between providers.

In this case, the total value of services is equal to the number of subscribers times the fixed service price. The total network value is equal to the value of services plus the value of connectivity, which is some percentage of the service value.

With zero market uncertainty, the services and content are not differentiated by feature set or content selection. Most users only value the service bundle as average, and thus they only value the transport connecting them to the service bundle as average. The bell curve in Figure 1 has been reduced to a single point in the middle.

High Market Uncertainty

When market uncertainty is high, the “walled garden” provider has no way to know which feature set will be seen by the users as best, or what content users are most interested in. The provider essentially has to guess. The value of this guess is random for each user, so the overall average for all users is the distribution average.

Once the provider has decided – however arbitrarily – which services to offer, it has a little easier time deciding

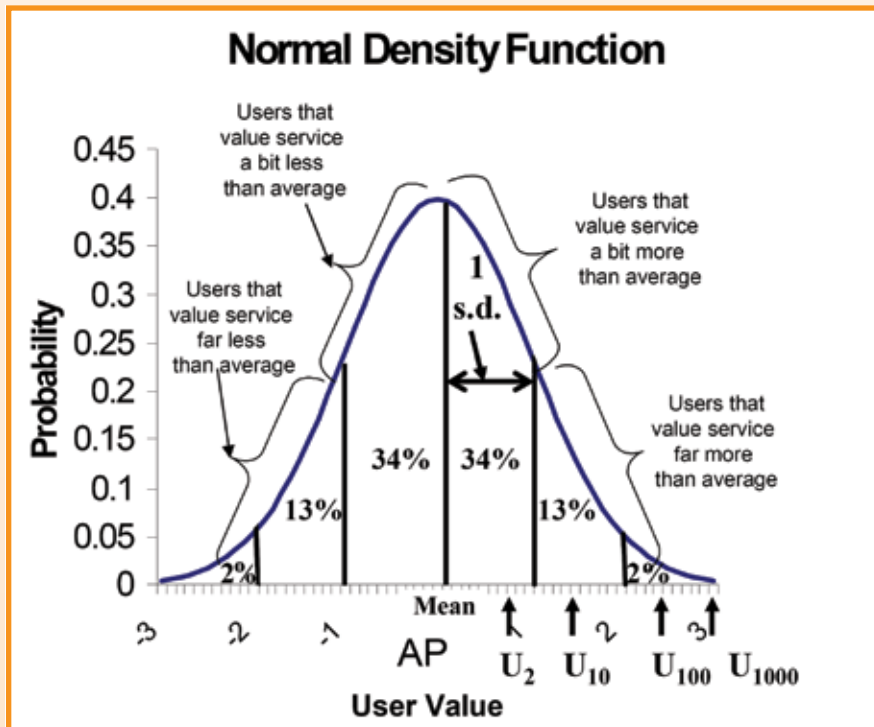


Figure 1. The horizontal axis represents how much a user would be willing to pay for a service. The central point represents the price that 50 percent of potential users would be willing to pay. Around the middle, within one standard deviation from the mean, are the 68 percent of users that are moderately happy with the service. The users on the far right are very happy with the service and value it far more than the average user does.

what to charge for the bundle. Market research and experience should tell it how many users will subscribe at any given price level. Since users value the service bundle differently, a few of them (like the top 2 percent shown in Figure 1) may be willing to pay very high prices for the bundle that the provider has chosen. But if the provider raises its price to what the top 2 percent will pay, it will drive away the other 98 percent. On the other hand, if it lowers its price to capture 100 percent of users, it will be undercharging nearly everyone else. (We're assuming that, in a condition of high market uncertainty, the provider doesn't have enough information to charge different prices to different users based on how much they want the services.)

Case 2: Network Neutrality

In a perfectly neutral network, users can pick services and content from many choices. The transport service provider offers transport services and, optionally, a bundle of network-based services that a user may or may not subscribe to. The transport provider keeps 100 percent of the connectivity revenues as well as any revenues from the services that the user selects. Because we believe that transport is more valuable to users when they value the available services more, one way to view this connectivity revenue is as a percentage of the service revenue (sort of a "network tax"). But if the user selects network-based services from other providers, these other service providers are paid the service portion of the total net-

To simplify the model, let's assume that the transport service provider gets almost all of the service revenues (in addition to its 100 percent of the connectivity revenues), simply because most users find it easier to deal with a single provider for both transport and network-based services, and they are happy with any bundle because of the lack of market uncertainty.

High Market Uncertainty

In a highly uncertain market, users will have different perceptions of their needs. In this case, the ability to choose becomes very important. Each potential user can select the best match from a variety of mass-market or niche network-based service or content providers.

There are "hard" and "soft" definitions of network neutrality. The "hard-line" definition for the Internet is that the network should have only one class of traffic: best effort. A "softer" approach would allow different tiers of traffic handling as long as quality of service (QoS) is sold in a fair and nondiscriminatory fashion to all users and service/content providers. We argue for soft network neutrality for several reasons.

It turns out that in many cases over a reasonable range of market uncertainty that a good strategy for the provider is to set the price right in the middle, where it will attract 50 percent of the potential subscribers. But this revenue stream will be much smaller than the revenues generated by the 100 percent subscription rate in a zero-uncertainty scenario. In fact, it will probably be about half of the zero-uncertainty revenue stream, if the mean price under uncertain conditions is (on average) about equal to the consensus price in the absence of market uncertainty.

Thus, a non-neutral network is more valuable when there is a low degree of market uncertainty and less valuable when there is high market uncertainty.

work revenue for those services. Providers or network-based services that guess better about what features users want can charge more for those services.

No Market Uncertainty

In a commodity market, the ability to choose doesn't really matter to users because they value all choices the same. As in the non-neutral case, 100 percent of potential users subscribe to one or another of the network-based services. Thus the total value of the network is the same as for the non-neutral network. However, in the neutral network the total network value is divided among the transport provider and the various providers of network-based services and content over the network.

Intuitively, it seems clear that when users want different features, choice is important to them. If there are many competing network-based service and content providers, we would expect the values of their offerings (from the user's point of view) to be randomly distributed. We assume for simplicity that this random variable describing the value of services to users is normally distributed. However, our model does not depend on a normal distribution, and most "reasonable" distributions illustrate the same results. As Figure 1 shows, about two thirds of the services will be seen as moderately good by the users – within one standard deviation of the mean – and a tiny percentage will be seen as truly great by the users. A user with 10 choices is likely to find one that is good but not "great;" the

High market uncertainty and feature-based competition allow service providers to charge more for more successful ideas, allowing the possibility of capturing a large potential market and winning big.

best match is most likely to be roughly 1.5 standard deviations from the mean (U10 in Figure 1). With 1,000 choices, almost any user can find an offering he considers "great" (U1000).

But even if there is only one choice (for a total of two choices for the users, which is illustrated by U2 in Figure 1) in addition to the service bundle offered by the network owner, the total value of the network will increase. Why? The revenue increase in neutral networks comes from two different sources. First is from the users who would have subscribed to the non-neutral network. With more choices available, some of these users will pick services they are willing to pay more for. They will all still have the same choice of bundled services that was available in the non-neutral network, but many of them will find services they value even more highly, and pay more for them.

Next, some of the users who would *not* have subscribed to the non-neutral network will now be able to find services they want. These potential new users were not willing to pay the price for the ISP's service bundle, but some of them will find competing network-based services to be a better match to their needs and will now subscribe to the neutral network in order to access those network-based services even if the price of the network is higher than the

pure connectivity cost in the non-neutral network.

The value of a neutral network depends on both the number of service/content choices and the shape of the "bell curve" showing the likely value of these choices. As users are presented with more choices, the more likely they are to find an ideal match. Remember,

low market uncertainty is represented by a tall and narrow bell curve, while high market uncertainty produces a low and wide bell curve. As the standard deviation of the distribution grows, so does the value of having more choices.

With low market uncertainty and many choices, the best service may not represent much of an improvement over



No Worries!
Control & monitor your SMATV headend from the office, on the road or from home....

AMT would like to introduce HEMS (Head End Management System)

- Monitor video & audio quality
- Reboot receivers
- Change channels on receivers
- Check signal levels for each receiver
- Auto scan to detect problems
- Email notification

Call 888.293.5856

AMT
advanced media technologies®

Advanced Media Technologies, Inc. - 720 South Powerline Rd., Suite G, Deerfield Beach, FL 33442
954.427.5711 - Fax 954.427.9688 - Toll Free 888.293.5856 - www.amt.com - sales@amt.com

Network neutrality boils down to providing users with an unbiased choice of services and content, which our model implies will promote greater user satisfaction as well as opportunities for innovators and entrepreneurs. What is this market uncertainty? In brief, it is the inability of service and content providers to predict what users will like and how users value the features of a service or the selection of content.

the “non-neutral” choice, because the choices are all so similar. However, with high market uncertainty, the choices are very different, and the user is much more likely to find a service that he or she greatly prefers to the non-neutral

choice. This means that a neutral network provider can charge more for the network connectivity than can a provider of a non-neutral network.

Thus, a neutral network is more valuable when there is a high degree of

market uncertainty and less valuable when there is low market uncertainty. Note that this is the opposite of the non-neutral network, which loses value when market uncertainty is high.

How Uncertain Is the Market?

There is reason to believe that market uncertainty today is quite high. The average price paid by users for a random bundle of services can be about \$100 in some markets. If market uncertainty were small, then few users would be likely to pay more than \$103. In actuality, however, many users pay as much as \$200 or more for service bundles including cable TV, e-mail, voice, and transport – implying that market uncertainty is very high. To illustrate this level of uncertainty consider the price one of the authors is paying. In the Boston area both Verizon and Comcast offer bundled Internet/email/voice/cableTV for \$99. Unfortunately, in Cambridge this “teaser rate” is not available. With two HD TVs both with DVRs and a mod-

KT Communication

We Share Your Vision - A Commitment to Quality and Integrity

KT Communication provides digital satellite video programming solutions for C-Band, Bulk, L-Band and Transport.



Our video platform delivers content to independent multi-channel video providers (MVP) that target hard-to-serve and alternative technology markets, including:

- Multi-dwelling and multi-tenant units, such as apartment complexes and duplexes, university residence halls and hospitals, vacation resorts, and real estate investment trusts with multiple units
- Private cable operators
- Wireless cable operators
- Franchise cable operators, particularly operators of small and rural systems



2409 N. Stadium, Columbia, MO 65202 • 877-485-3557 • 573-446-3693 • Fax: 573-446-9054
www.ktcom.tv

erate package of premium channels the author is paying \$167.50.

With high market uncertainty and 100 different choices, the best of the 100 choices is roughly 2.5 standard deviations from the mean. This would imply that the benefit from network neutrality is huge for both users and independent service and content providers.

Network neutrality may even benefit network providers. If users value connectivity on the basis of the services they can access, and if network neutrality allows them to access more valuable services, then the connectivity itself will be more valuable to them – and 100 percent of the connectivity revenues go to the network provider.

With network neutrality, users have services they value more, the transport service provider can make more money,

A single set of services from a single service provider is unlikely to match users' wants as well as a group of independent service providers would. A single transport service provider will never develop or select as many innovative services as a group of heterogeneous and independent providers are able to do.

around to promote investment in network infrastructure and research.

Conclusion

Our model illustrates the value of unbiased choices in neutral networks. It

of a neutral network be split among transport and other service and content providers so as to insure that they have an incentive to build infrastructure and create the varied service and content that meets uncertain user needs?" **BBP**

References

1. Gaynor, M., Bradner, S., Statistical Framework to Value Network Neutrality, Media Law & Policy, New York Law School, Accepted for publication March 2008.
2. Yuksel, M, K, Ramakrishnan, S. Kalyanaraman, J. D. Houle, and R. Sadhvani, Value of Supporting Class-of-Service in IP Backbones, Proceedings of IEEE International Workshop on Quality of Service (IWQoS), Chicago, IL, June 2007.

About the Author

Mark Gaynor is Assistant Professor at the Boston University School of Management, Information Systems Department. He holds a Ph.D. in computer science from Harvard University, and his research focuses on Internet-related topics. His recent book based on his PhD thesis is Network Service Investment Guild: Maximizing ROI in Uncertainty Markets, published by Wiley in January 2003. He can be reached at 617-353-4159 or at mgaynor@bu.edu.

Scott Bradner is the University Technology Security Officer at Harvard University. He has been working on the technology of the Internet since before it was called the Internet. In particular he has been very active in the Internet Engineering Task Force (IETF). He also writes a weekly column for Network World. He can be reached at 617-495-3864 or sob@sobco.com.

Given a pie that could be four times as big for a neutral network, a reasonable set of regulations, and a market where users can select what matches their needs best, there will be an incentive to build network infrastructure along with a fertile environment conducive to the independent service and content providers, creating a surplus benefit to society.

and the competitive environment for network services and content is good for innovation and investment. The much greater value of a neutral network implies that with a very high likelihood there exists a "fair" distribution of total revenue to split between infrastructure and content service providers.

Depending on the parameters of our model the infrastructure provider can make as much as before while still leaving this amount on the table to split between other content providers and future network infrastructure. If market uncertainty is as large as it appears, then there is plenty of money to spread

demonstrates that given a set of "reasonable" assumptions, it is highly likely that the total revenue pie is much bigger than in non-neutral networks. Given a pie that could be four times as big for a neutral network, a reasonable set of regulations, and a market where users can select what matches their needs best, there will be an incentive to build network infrastructure along with a fertile environment conducive to the independent service and content providers, creating a surplus benefit to society.

The important question is no longer, "Should the network be neutral?" Rather, it is, "How can the bigger pie