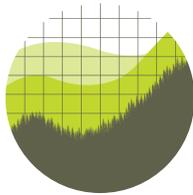


Internet Benefits:

Consumer Surplus and Net Neutrality



Institute for
Policy Integrity

new york university school of law

J. Scott Holladay

A.J. Glusman

Steven Soloway

Policy Brief No. 10

October 2011

Across the globe, many of the defining political and cultural moments of the last ten years centered around the electronic networks that connect us: political protests have been organized from smart phones and news of natural disasters has broken online and spread virally.

Naturally, as the number of things we can do on the Internet increases, the amount of time we spend online has grown.

Most people do not consider how the combination of content and communications infrastructure improves their lives. Part of the beauty of the web is that end-users who use the internet recreationally do not have to think about where the entertainment is coming from or how it got to their home. You can just open a laptop, charge up your smartphone, or turn on Web TV and instantly connect to a rich universe of information.

While end-user access has been revolutionized over the last decade, the backend network that delivers content has been changing as well. Behind the scenes, new technologies have had far-reaching consequences for people who use the Internet and the companies that cater to them.

But not all technological advances have been beneficial. One, which has allowed Internet Service Providers (“ISPs”) to change how they deliver content to consumers, has threatened “network neutrality,” the status quo in which all content that travels across the Internet is treated equally.

This policy brief describes how a weakening of the principle of network neutrality might impact the Web. Based on an analysis of Internet usage, it finds that Internet infrastructure and content work together to generate huge economic benefits for consumers—possibly as much as \$5,686 per user, per year.

Eliminating network neutrality, as some have proposed, may reduce incentives to invest in Internet content and infrastructure. The Federal Communication Commission’s Open Internet Rule was announced December 23, 2010. While imperfect, it provides a layer of protection for important portions of the Internet by prohibiting ISPs from discriminating against content-providers based on content or their ability to pay for access to end-users. This could help ensure that the Internet’s value is not threatened by the misaligned incentives of some market actors.

What exactly is “the Internet”?

The Internet consists of two interconnected goods: the infrastructure of wires and hardware that connect us to the web and the content and applications that travel across those wires. Internet infrastructure is a network of wires and hardware that brings content into homes and businesses. The content that is provided over those wires runs a long gamut: from text web sites, to calls on Skype, to streaming movies on Netflix.

Economists refer to goods that are traditionally or necessarily consumed together as complimentary goods. Pairs of goods like peanut butter and jelly, DVDs and DVD players or iPhones and applications are inexorably linked. Changes in the price or quality of one good can have an impact on its counterpart.

Internet infrastructure and content make a pair of complimentary goods. An increase in the price of infrastructure should lead to a reduction in demand for Internet content and *vice versa*. Improved Internet content will lead to an increase in demand for Internet infrastructure. It is a self-reinforcing cycle. As Internet infrastructure improves, the quality of Internet content increases to take advantage of faster access to consumers. As Internet content becomes more valuable (and more bandwidth intensive) Internet infrastructure expands to handle the load.

The value that consumers derive from the Internet is dependent on both the quality of the infrastructure and content.¹ If the infrastructure cannot handle the demands of users for high bandwidth content provided at fast speeds, then consumers will derive less value from the Internet than they would from a faster Internet. At the same time, a high-speed connection is useless if consumers are not interested in the content that they receive quickly. If either Internet content or infrastructure lags, the consumer experience will suffer. The future of the Internet cannot be evaluated without considering how content and infrastructure must evolve together.

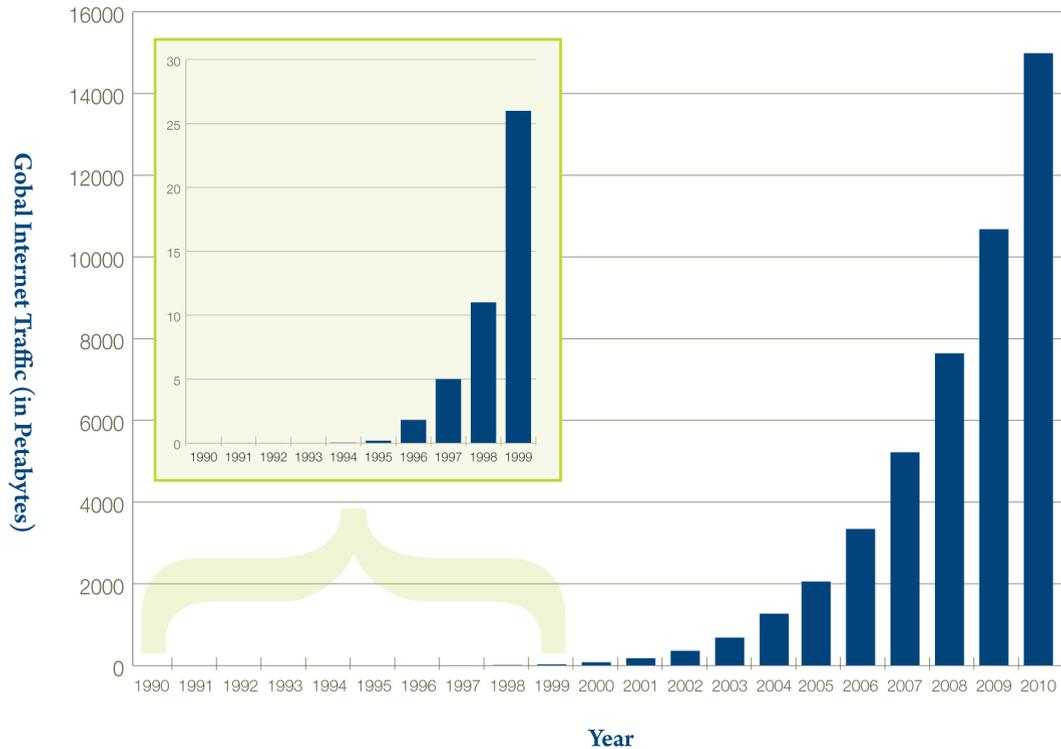
ISPs’ capacity to keep up with increased Internet use is critical because Internet traffic has been growing at an exponential rate. In the late 1990’s, Internet traffic doubled in size every 90 days.² Growth has slowed slightly in the 2000’s, but still the amount of data moving over the Internet is estimated to have doubled every two years.³ In 1990, Internet traffic in the United States was about 1 terabyte a month,⁴ today it takes only 25 Netflix users to generate that the same amount of traffic.⁵ Strong growth is expected to continue for the foreseeable future: Cisco projects that Internet traffic will more than double between 2012 and 2015.⁶

Recently, it became possible for Internet infrastructure owners to determine the source, destination and content type of every packet of data that moves across their network. Using a technology known as “Deep Packet Inspection” (“DPI”), ISPs can tell not only if a packet contains information that is part of an email, webpage or an Internet video, but also if it came from Amazon, Ebay or some other website. This represents a huge advance in ISPs’ insight into the data that is moving across their networks, but has potentially problematic applications that will be explored more fully below.

Consumer Surplus Generated by the Internet

How much would you be willing to pay for the Internet? Many people would be willing to pay more than they currently do. The difference between what users would be willing to pay and what they actually pay is what economists call “consumer surplus.” This surplus is a measure of how much better off the Internet makes consumers, and is one way to measure the benefits a market creates for consumers.

Estimated Internet Growth Over Time



Source: Cisco Visual Networking Index: Forecast and Methodology, 2010-2015

Think of it like this: the amount you are willing to pay for Internet infrastructure and content bundled together is the total value of the Internet to you. The amount that you pay your ISPs and content providers is subtracted from that value and the remainder is the value that you get to keep.

Consumer surplus is not the only benefit that the Internet generates. ISPs and content developers are each better off because they sell to consumers for more than the rock bottom price they would be willing to accept. This analog to consumer surplus is called producer surplus. Producer surplus is another component of the benefits generated by the market. Rather than going to consumers, these benefits are typically expected to go to the shareholders of the ISPs and Internet content companies.

The consumer surplus and producer surplus combined are equal to the total benefit that the Internet generates. The more surplus a market generates, the better it is for consumers in that market. Economists typically worry more about the amount of the surplus and less about the distribution of that surplus. So, from a purely economic perspective, if a market generates billions of dollars in surplus, but the entire surplus goes to producers, this does not necessarily mean that something is wrong with the market.

Similarly, different types of companies might be fighting for a piece of the same producer surplus. In the Internet market, ISPs and content companies are battling over the spoils of the Internet market, but as long as the total amount of surplus is as large as possible, who wins that struggle might not be that important.

So how do we calculate the value of consumer surplus generated by the Internet? One way would be to ask people how much they would be willing to pay for the Internet and compare that to how

much they spend on Internet service and content. The difference would be their consumer surplus. This might not be the best approach however; most people do not have a good sense of how much they would be willing to pay for the Internet or even how much they actually pay. Eliciting the true valuation is therefore difficult, and in many cases the answer is that consumers just do not know how much value they are receiving.

Fortunately, there is another approach. Economists have some tricks for figuring out how much value consumer products like the Internet generate. A pair of economists from the University of Chicago and Stanford University developed a technique for valuing the Internet (and other consumer products) by looking at the amount of time consumers spend using them.⁷ The logic is simple: the more time you spend on the Internet, the more you must value it. Consumers have a myriad of ways to spend their free time and the choices that they make should be based on what they enjoy doing the most.

So if consumers spend more time on the Internet and less time watching television or reading, then the Internet must bring them more surplus. The trick is figuring out the value of that time. Professors Austan Goolsbee and Peter Klenow use income to value time. Think of it like an hourly worker who can always pick up more hours. If you spend an hour on the Internet, that must be worth at least as much as you could have earned by working during that time. That basic insight lets us translate time spent on the Internet (and annual income) into estimates of consumer surplus.

Survey data on how much time people spend on the Internet is available from the Pew Research Center.⁸ It has conducted occasional surveys asking how much time people spend online each day and what they do online.⁹ The surveys also ask a series of socioeconomic questions such as age, sex and, crucially for this analysis, income. When we take the data on how people spend their time and how much money they make and run it through the formula that relates time and income to consumer surplus, we get an idea of how much consumers benefit from the Internet.

The results suggest that the consumer surplus generated by the Internet is very large.¹⁰ The average survey respondent spent 114.5 minutes a day on the Internet recreationally. The benefits that use generates are equivalent to 5.2% to 7.1% of income.¹¹ If we use the median income value of Pew's survey, we find that individual consumers collect between \$4,155 and \$5,686 worth of value from the Internet per year. This estimate is big, but it is in the same neighborhood as those found by Goolsbee and Klenow. They found that the consumer benefits of the Internet were somewhere between 2% and 3% of total income.¹² The amount of time consumers spend on the Internet suggests that they receive a great deal of benefit from access.

Changes to the way the Internet is managed could increase these benefits even further or reduce the amount of surplus the market generates. For that reason, any policy changes need to be carefully considered and their impact sufficiently understood. Poorly made Internet policy could put a lot of consumer benefits at risk.

This analysis suggests that demand for the Internet is likely to be inelastic,¹² which means price increases are met with less than proportional reductions in demand. For example, if the price of the Internet went up by 10%, the demand for Internet would drop by less than 10%. Inelastic demand for the Internet implies that end-users will not respond strongly to changes in price for the bundle of Internet content and infrastructure.¹³ If the price ISPs charge increases rapid-

If we use the median income value of Pew's survey, we find that individual consumers collect between \$4,155 and \$5,686 worth of value from the Internet per year.

ly, most consumers will keep their subscription and pay the higher price. Inelastic demand also means that the Internet market generates large benefits for consumers. Our analysis suggests that these benefits are worth thousands of dollars a year and are probably growing over time.¹⁴

Network Neutrality: Internet Content vs. Infrastructure

Network neutrality refers generally to treating equally all the information that flows across the Internet no matter the source, destination or type of content. The Internet has evolved with an ‘end-to-end’ architecture, which means that when an Internet user places a packet of information onto the network, a series of intermediaries move it from its source to its destination without speeding or slowing its journey.

The distinction between Internet content and Internet infrastructure plays a crucial role in the network neutrality debate. Internet content providers develop web sites and applications and then pay to upload them onto the Internet. End-users then connect to those sites and apps and download the data onto their desktops or laptop computers, phones or tablets, which are connected to the Internet by an ISP. Content companies do not pay for the traffic that the downloading of their information generates and have no control over how it reaches end-users. The interaction between content providers and ISPs has a big impact on the quality of Internet that consumers enjoy, but there is little to no coordination between the two industries.

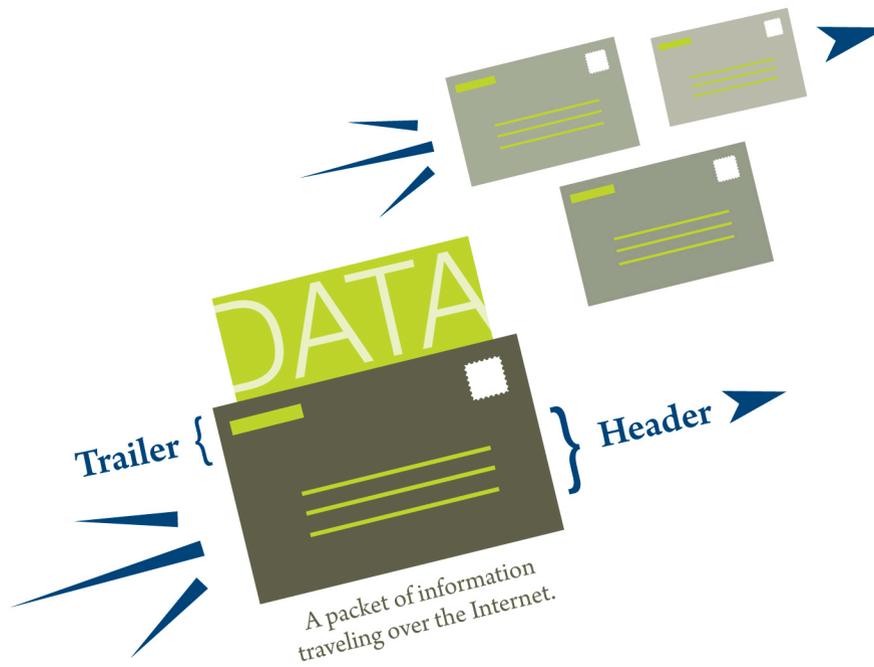
Advances in technology, such as DPI, have changed the existing dynamic of how information moves around the Internet. Internet infrastructure owners can use DPI in a number of ways. They can “prioritize” packets of streaming video by delivering them immediately while forcing other types of content to wait. Prioritized video service loads faster and is less likely to stutter or experience delays when the ISP’s network is crowded. ISP’s can also use DPI to prioritize packets from a particular site like Amazon and/or slow packets from competing sites like Ebay. The advent of DPI has given ISPs the unprecedented ability to shape the flow of traffic across their networks.¹⁵

Before the advent of DPI, ISPs didn’t know which content developers’ packets were flowing across their networks. Now that ISPs have this knowledge, they have the ability to slow or even block packets from a specific content provider. This could allow ISPs to take advantage of their position between end-users and content. For example, an ISP could request a fee for priority access to end-users from a search provider. If a search provider paid for it, the ISP could deliver that search provider’s results directly while slowing the delivery of its competitor’s results.¹⁶ Over time, this would drive customers to the faster search engine, regardless of the quality of their results.

The ability of ISPs to speed access to end-users is limited by available technology and physics. You can think of packets arriving at an ISP as similar to people waiting in line at a store. If there are lots of cashiers working, no lines will form. If the store gets busy or does not have enough cashiers, however, a line forms, and the first person in line is the next one to check out. Before DPI, the policy for information packets was first come, first served. But, in this analogy, DPI allows the store to identify all the people in line, pluck the VIPs out, and move them directly to a cashier while everyone else in line waits. If there is no line, prioritization does not make a difference, but when the line forms, prioritization allows selected packets to jump the line.

Network neutrality has the potential to affect the distribution of producer surplus generated by the Internet. If content developers are able to pay ISP’s for priority access to consumers, it allows ISPs to capture a portion of the surplus that used to go to content creators.¹⁷ This allows for the possibility to shift investment out of content and into Internet Service Provision. No one knows for sure how this shift would affect the Internet, but there is a risk that ending network neutrality could reduce investment in both Internet content and Internet infrastructure.

An Illustration of Deep Packet Inspection



Eliminating net neutrality may have an impact on the market by inducing ISPs to create network congestion. That is, if there is no line at the cash register, then there is no benefit to being a VIP (who has the option of jumping the line). If ISPs invest a lot in infrastructure and eliminate congestion, then no one will pay for prioritization. In the absence of network neutrality, ISPs would have conflicting incentives regarding investment in congestion reducing infrastructure improvements.

On the one hand, those investments would make an ISP more attractive to customers, but on the other hand, they would make the VIP status offered to some content providers less valuable.¹⁸ Without net neutrality, it is more likely that ISPs would opt not to invest. This would make their networks more congested, which in turn would allow them to charge content companies for the right to “jump the line.” If ISPs were to invest in infrastructure, congestion would be reduced, and they would not be able to charge as much for prioritization.

Net neutrality protects incentives for content providers to invest. If ISPs are able to extract more relative surplus from content developers, this might discourage potential start-ups from testing the market. Such a fear would not be unfounded – if the start-up cannot afford to pay for initial prioritization, they may have trouble competing.

They may also fear that, once they become profitable, a large share of their revenue will be extracted by ISPs. Between content providers fearing the loss of a share of their revenue to ISPs and the ISPs investing less in infrastructure, in a non-net-neutral future, consumer benefits could take a big hit. Some ISPs have claimed that ending network neutrality will allow them to develop new types of content, such as 3D video or more advanced cloud computing services.

This argument is one of the keys to the network neutrality debate. If the status quo is changed and ISPs are allowed to prioritize content, they will likely earn more money, but there is no guarantee that they will invest that additional revenue in their networks. Much of that additional revenue will reward past investments, rather than incent the development of new infrastructure. If they return the additional revenue to shareholders, or use it to invest in their own Internet content, then the future benefits from the Internet are at risk.

Conclusion

The Internet consists of two complimentary goods: Internet service and Internet content. These two industries must work in concert to provide consumers with what we think of as a single integrated good: The Internet. The interactions between these markets have evolved with the Internet and any technological or regulatory change in one of the markets will have an impact on the other. Economics has developed tools for predicting how these changes will affect each market and the consumers that participate in them.

The Internet has evolved from an esoteric service to connect university computing labs into the primary communication tool of our age. This evolution has been associated with huge increases in benefits to consumers. Our estimates suggest that recreational use of the Internet could be worth as much as \$5,686 in economic benefits per user per year. There is evidence that the size of this benefit has been increasing over time. If this is true, increasing recreational Internet use has gone from providing no benefit to providing benefits worth thousands of dollars per person per year in less than twenty years.

The Internet currently operates under a *de facto* network neutral regime as it has since it began. This regime has generated huge amounts of consumer surplus and the amount of surplus is growing over time. No one knows exactly what will happen if we end this neutral regime and allow ISPs to prioritize content, but the burden of proof should be on the ISPs that want to change the status quo. Those who argue that we should change the regime should be called upon to show that these changes will not deplete the huge surplus that the Internet generates.

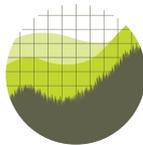
On October 22, 2009, the FCC formally proposed a net neutrality regulation, entitled “Preserving the Open Internet.”¹⁹ The rule includes a provision that prohibits ISPs from “discriminating.” That provision states, “[s]ubject to reasonable network management, a provider of broadband Internet access must treat lawful content, application, and services in a nondiscriminatory manner.”²⁰ This nondiscrimination rule would prohibit a broadband ISP from discriminating against, or in favor of, any content, application, or service, subject to “reasonable network management.”²¹ The FCC noted that the nondiscrimination rule prevents an ISP from “charg[ing] a content, application, or service provider for enhanced or prioritized access to the subscribers of the broadband ISP.”²² While the rule does not extend the non-discrimination provision to wireless Internet access, and includes exemptions for “managed services” that will need to be tightly monitored to ensure they do not become overly broad, it does provide an important measure of protection for net neutrality.

Without these protections, there is a real risk that net neutrality will be threatened and ISPs will be encouraged to reduce investment in infrastructure, which will increase congestion and prioritization revenues. If Internet infrastructure is allowed to languish, then the incentives to invest in content will be reduced, and the benefits that consumers receive from the Internet could shrink. Network neutrality opponents face a high burden of proof to convince regulators that ending neutrality will protect the existing benefits of the Internet.

Notes

- 1 Although there are similarities between the analysis of consumer uses of the Internet and the productivity of B2B type applications on the Internet, this Policy Brief only considers the consumer portion of Internet satisfaction.
- 2 REED HUNDT, YOU SAY YOU WANT A REVOLUTION 224 (2000).
- 3 MINNESOTA INTERNET TRAFFIC STUDIES FOR 2002-2009, <http://www.dtc.umn.edu/mints/>; Andrew M. Odlyzko, Internet traffic growth: Sources and Implications, in *Optical Transmission Systems and Equipment for WDM Networking II*, Table 1 (Benjamin B. Dingel ed., 2003).
- 4 U.S Internet traffic is measured by total “backbone” traffic, which refers to the data routes between all “high- capacity network centers,” and which includes all commercial, academic, and corporate traffic.
- 5 SANDVINE, GLOBAL INTERNET PHENOMENA SPOTLIGHT: NETFLIX RISING 5 (May 2011), http://www.sandvine.com/downloads/documents/05-17-2011_phenomena/Sandvine%20Global%20Internet%20Phenomena%20Spotlight%20-%20Netflix%20Rising.pdf and author’s calculations.
- 6 CISCO, CISCO VISUAL NETWORKING INDEX: FORECAST AND METHODOLOGY , 2010-2015, Table 9 (2011), http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html.
- 7 This cost estimate represents the per-subscriber fixed costs of broadband Internet. We assume that the short run marginal costs to ISPs (the cost of adding one more user to the network) are negligible. See Seung-Jae Shin, Hector Correa & Martin B.H. Weiss, A Game Theoretic Modeling and Analysis for Internet Access Market, (Center for Technology, Policy, and Industrial Development, Program on Internet and Telecoms Convergence), Jul. 22, 2002, at 13, http://dspace.mit.edu/bitstream/handle/1721.1/1481/shin_correa_weiss.pdf?sequence=2. The average amount of money spent on Internet content is based on a weighted average of survey respondents from a Pew Internet Survey collected in October of 2010, which is available at [http://pewinternet.org/Shared-Content/Data-Sets/2010/November-2010--Paid-Content-\(Omnibus\).aspx](http://pewinternet.org/Shared-Content/Data-Sets/2010/November-2010--Paid-Content-(Omnibus).aspx).
- 8 Austan Goolsbee & Peter J. Klenow, Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet, 108-113 (Nat’l Bureau of Econ. Research , Working Paper No. W11995, 2006).
- 9 Pew Internet and American Life Project, www.pewinternet.com.
- 10 Pew Internet & American Life Project, Personal Networks and Community Survey, August 22, 2008, <http://www.pewinternet.org/~media/Files/Questionnaire/2009/PIAL%20Networks%20FINAL%20Topline.pdf>.
- 11 See Goolsbee and Klenow for some important caveats to this calculation. This approach values the benefits of leisure time using the wage rate, which may be inappropriate. If a survey respondent values their leisure at less than their wage rate (for example, if they would like to work more, but their hours are capped) then this calculation will overstate consumer surplus. Even if we assume that consumers value leisure at only 50% of their wage rate, the consumer surplus associated with recreational Internet use is between 2.6% and 3.55% of “full” income.
- 12 This figure represents consumer surplus calculated as the area under a linearized demand curve. The area under the curve (as a share of full income) is $0.5 * LI/\eta$, where LI represents the fraction of leisure time respondents spend online and η represents the elasticity of demand for leisure time. H is equal to $\sigma(1-LI(1-FI/W))$, where σ is the elasticity of demand between Internet and all other goods and services, FI is the cost of Internet access and W is the full wage rate. We calculate LI and W from the Pew survey data and take FI from the Goolsbee and Klenow paper. Using the σ calculated in the Goolsbee paper produces a consumer surplus of 5.2% and using σ calculated from the Pew data produces a consumer surplus of 7.1%.
- 13 Goolsbee and Klenow use a different data set that is larger, but also older. The Forrester data they used was collected in 2005, but consists of over 60,000 respondents. The Pew data used for these updated calculations was collected in 2008, but has only about 1,200 responses. For that reason the Goolsbee and Klenow calculations are more precisely estimated, but in a fast moving industry like the Internet, it is possible that these consumer surplus figures change quickly. The survey questions are also worded in slightly different ways. Forrester asks respondents how much time they spent on the Internet for personal use, whereas the Pew question is more open ended, asking only how much time the respondent spent online without specifying for work or personal reasons. The model is flexible enough to use the responses to both questions, assuming that all of a consumer’s (waking) time could be spent at work. If this assumption does not hold and workers are not able to add additional wages by working additional hours up to 16 hours a day then the consumer surplus numbers are not directly comparable.

- 14 Using the Goolsbee and Klenow approach on the Pew data yields a σ of 0.95. Any σ less than 1 represents inelastic demand. Note that Goolsbee and Klenow estimate σ to be in the 1.3 to 2 range. Both estimates are consistent with the idea that consumers are not particularly sensitive to price changes for Internet content and access.
- 15 Examples of other goods with relatively inelastic demand are cigarettes and the short term demand for gasoline.
- 16 The evidence that consumer surplus generated by the Internet is growing comes from comparing our estimates to Goolsbee's and Klenow. It appears that consumer surplus as a fraction of income has nearly doubled between 2005 and 2008. This result should be interpreted with care; we used different data sources so the results are not directly comparable.
- 17 For further reading on the history and impacts of Deep Packet Inspection see Nate Anderson, *Deep Packet Inspection Meets 'Net Neutrality*, *CALEA, ARS TECHNICA*, JUL. 25, 2007, [HTTP://ARSTECHNICA.COM/HARDWARE/NEWS/2007/07/DEEP-PACKET-INSPECTION-MEETS-NET-NEUTRALITY.ARS/](http://ARSTECHNICA.COM/HARDWARE/NEWS/2007/07/DEEP-PACKET-INSPECTION-MEETS-NET-NEUTRALITY.ARS/).
- 18 If individual search providers take steps to enter into an exclusive arrangement that forecloses competition by other sites, they may run afoul of antitrust laws. See *United States v. Microsoft*, 253 F.3d 34, 67-70 (D.C. Cir. 2001).
- 19 It is likely that the competitive content market would prevent content providers from passing on this cost to consumers. See Inimai M. Chattiar & J. Scott Holladay, *Free To Invest: The Economic Benefits of Preserving Net Neutrality* 28 (Institute for Policy Integrity, Report No. 4, 2010), available at, <http://policyintegrity.org/publications/detail/free-to-invest/>.
- 20 For a theoretical description of this phenomenon, see Shin, et al., *supra* note 10.
- 21 Preserving the Open Internet, Broadband Industry Practices, 74 Fed. Reg. 62,638, 62,638 (proposed Nov. 30, 2009) (to be codified at 47 C.F.R. pt. 8).
- 22 *Id.* at 62,646-47.
- 23 *Id.* at 62,661.
- 24 *Id.* at 62,646.



Institute *for*
Policy Integrity
new york university school of law