

Subsidizing Creativity through Network Design: Zero-Pricing and Net Neutrality

Robin S. Lee and Tim Wu

The “net neutrality” debate, as it has emerged over the last five years, is a social, political and economic debate over the public information network known as the Internet and the duties of its private carriers, which include telephone and cable companies and other Internet service providers (ISPs). In the early 2000s, questions surrounding the rights of Internet carriers to block certain network attachments and control access to emergent applications or content providers led to a call to protect “network neutrality” (Wu, 2003).¹ The debate raises familiar questions for students of travel or communications infrastructure: for hundreds of years, courts and governments have struggled over the duties that carriers like ferries, railroads, or telephone companies owe to the public by virtue of their necessity to economic and social welfare. While the net neutrality debate has many aspects, in this paper we focus on one crucial issue: the de facto rule

¹ In March 2005, the Federal Communications Commission fined an Internet service provider, Madison River, for blocking an Internet telephony service. In December 2006, the FCC imposed net neutrality rules on AT&T as a condition of its merger with Bell South, forcing it to commit itself “not to provide or to sell to Internet content, application, or service providers, including those affiliated with AT&T/BellSouth, any service that privileges, degrades or prioritizes any packet transmitted over AT&T/BellSouth’s wireline broadband Internet access service based on its source, ownership or destination” (Reardon, 2006). In July 2008, the FCC criticized cable firm Comcast for deliberately interfering with transmissions occurring through BitTorrent peer-to-peer software, which allows large and/or popular files to be downloaded simultaneously by many users, and required that such interference cease.

■ *Robin S. Lee is Assistant Professor of Economics, Stern School of Business, New York University, New York City, New York. At the time this paper was written, he was Research Scientist, Yahoo! Research (Microeconomics), New York City, New York. Tim Wu is Professor of Law, Columbia University Law School, New York City, New York, and a Fellow at the New America Foundation, Washington, D.C. Their e-mail addresses are <rslee@stern.nyu.edu> and <wu@pobox.com>.*

prohibiting consumers' ISPs from charging fees to content providers for access to their customer base.

The Internet can be understood as an information network that intermediates between different groups of agents comprising *users* and *content providers*, with the critical understanding that users can also act as content providers (for the purposes of discussion, we will use the term "content" loosely, referring to all types of media, applications, retailers, and services available online). Since the Internet's spread through academia in the 1980s and mass popularity in the 1990s, it has maintained a pricing structure that is unique among information networks: users and content providers typically pay ISPs *access* fees—fixed fees to get on the Internet at all—and *usage* fees—variable fees paid based on time or bandwidth usage; however, there have not generally been any additional charges for one user of the network to reach another user or content provider. For example, content providers such as Google and Wikipedia, while paying for their own Internet access and usage, do not directly pay the ISPs of users they reach.

This began to change in the mid-2000s as certain Internet service providers began to discuss their desire to charge certain users of the Internet—notably, large application or content providers—additional fees to reach their subscribers. For example, AT&T, an ISP, might demand that content providers such as Google and Wikipedia pay AT&T to access AT&T's customers; failure to comply would result in AT&T blocking traffic from those sites to its customers (and also preventing its own customers from reaching those sites). Following usage in the telephone system, we refer to these fees as *termination* fees: AT&T would charge content providers a fee to deliver their packets much like it charges other telephone networks a fee to "terminate" their calls.

Today, through historical practice, there exists a *de facto* ban on termination fees—also referred to as a "zero-price" rule (Hemphill, 2008)—which forbids an Internet service provider from charging an additional fee to a content provider who wishes to reach that ISP's customers. The question is whether this zero-pricing structure should be preserved, or whether carriers should be allowed to charge termination fees and engage in other practices that have the effect of requiring payment to reach users. This paper begins with a defense of the *de facto* zero-price rule currently in existence. We point out that the Internet, as an intermediary between users and content providers, exhibits pricing dynamics similar to other intermediaries in "two-sided markets." In particular, we posit that the Internet's absence of payments from content creators to users' ISPs facilitates the entry of content creators. In that respect, the rule provides an alternative implementation of the policy goals provided by the intellectual property system and achieves functions similar to copyright and patent law. The rule also helps avoid the problems of Internet fragmentation, in which content providers who do not reach agreements with ISPs cannot access all customers, and consumers on a single ISP are foreclosed from accessing their content.

We then consider some of the main arguments against restricting termination fees. Although we concede that limited instances may justify certain deviations from this norm, ultimately we argue that a zero-price rule has helped make the Internet distinctive from other networks in terms of its level of creativity and social usefulness.

Zero-Pricing and Net Neutrality

The general practices specified by the phrase “network neutrality” emerged not as a closely considered policy decision, but as a consequence of how the Internet was designed and how it spread. We claim, however, that an economic case for the pricing rules inherent in a neutral network can be found in the theory of two-sided markets. This theory suggests the de facto ban on termination fees may be interpreted as a policy that provides a subsidy to content creation and provision. For a number of reasons, this subsidy appears to have been one of the forces generating the enormous wave of innovation in services and information in the last few years as well as spawning new forms of creative activity.

Internet Pricing Contrasted with Telephone and Cable Television Networks

Currently, Internet users pay access and usage fees to their service provider and then can reach any other user who is similarly connected to the Internet. The overall network does not, by its own design, distinguish between content providers and users. Consequently, content providers—who may also be users—are also able to reach an audience consisting of every single Internet user. These norms and expectations, which have created a de facto ban on termination fees, stands in sharp contrast to what is standard practice on other important information networks, like the telephone and cable networks.

One reason for the differences between networks is rooted in history. The Internet was conceived by various visionaries, particularly the Department of Defense researchers J. C. R. Licklider and Robert W. Taylor, as a “network of networks” or an “intergalactic network” that would make it possible for users of any single computer network to reach users on any other network (Licklider and Taylor, 1968). In its original, noncommercial setting, fees to access the network were paid by universities, government, and research departments. There was no motive to charge termination fees to content providers. Government regulations designating the Internet a noncommercial network also discouraged any such fees. On the technological side, Internet protocols were designed to create a network that could be universally used by different parties with very little centralized knowledge as to who was reaching whom; as a result there was little power to track or bill for reaching certain entities. In the early years, the complexity of trying to incorporate billing capabilities might have doomed the project.

In contrast, telephone and cable television networks were designed from the outset as commercial networks, where payments were the prerequisite to connectivity at all. The Bell telephone system was, from its origins, extremely careful about who would be allowed to connect to Bell customers. In the two historically “competitive” periods in wire-line telephony—from the 1890s to the 1910s and from the 1980s through the early 2000s—calling a Bell customer generally meant paying a termination fee to the owner of the local switch. Since the 1970s, termination fees on the telephone system have been regulated based on fears that the Bell companies would use their “termination monopoly”—their exclusive access to

customers—to charge exorbitant rates and bankrupt any would-be competitor. Cable television networks are priced in a similar, but in some ways opposite, manner. From the beginning, cable companies were faced with demands for fees to access the content they needed from broadcasters, copyright owners, or other sources of content. As opposed to charging to reach their users, cable companies currently pay content providers for the right to carry content on their cable networks.

The Consequences of Pricing Decisions in Two-Sided Markets

The historical reasons why the Internet has developed without termination fees does not address the issue of whether the principle is economically justified. One potentially powerful justification emerges from the economic theory of two-sided markets.

Generally, the Internet, as a network (or network of networks), can be seen as a market or intermediary that facilitates the interaction of two main groups: users and content providers. Other well-studied examples of two-sided markets include payment systems, such as credit cards or online services; hardware-software markets like videogames or operating systems; retail marketplaces such as bazaars, shopping malls, or auction houses; matching markets, such as nightclubs or job sites; and advertising exchanges, such as online advertising platforms as well as commercial telephone books.

One prevalent aspect of two-sided markets is the presence of “network effects” (Farrell and Saloner, 1985; Katz and Shapiro, 1986; Liebowitz and Margolis, 1994), where the value of a service generally increases in the number of users. Most obviously, the network is more useful to users the more people there are to e-mail, videoconference, or instant message. There are also indirect network effects in that users benefit from more content providers, and these providers also benefit from having more users to reach. “Negative” network effects are also possible: for example, in networks with congestion, the value of using a network may decrease with additional users.

The recent literature on two-sided markets has studied the pricing decisions of a network provider or platform intermediary (Rochet and Tirole, 2003, 2006; Caillaud and Julien, 2003; Armstrong, 2006; Hagiu, 2006) in the presence of such network effects and has shown that in settings with transactions costs or limits on side-payments that are allowable between agents, the division of prices charged between sides of the market matters greatly. When there are such frictions, charging fees to one side as opposed to the other—even if the sum of prices across sides is the same—can affect who uses the network, overall transaction volume, and ultimately the efficiency of the market.

Theory predicts that the optimal pricing decision for a social planner as well as for a monopolist platform provider might involve “subsidizing” one side of the market over the other(s)—that is, possibly charging one side below the marginal cost of providing service. A canonical example of this phenomena are credit card companies. Credit cards companies do not charge both sides of the market—merchants and consumers—equally. Rather, they charge a transaction fee to

merchants and typically reward consumers for using their cards with frequent flier miles or other benefits. Because using a credit card is often not only “free” for a consumer, but in fact subsidized, the effect is to encourage consumers to use credit cards more than they would if fees were charged to both sides of the market. The division of pricing matters in this case partially because merchants are often prohibited from charging consumers different prices depending on whether the method of payment is by cash or credit; in other words, since merchants are unable to “unwind” these fees, pricing has a “nonneutral” effect.

If one side is subsidized at the optimum, theory predicts that the side receiving the subsidy will be the side that either has a higher elasticity of demand with respect to price, or the side that exhibits stronger cross-side network effects—that is, an additional agent on one side of the market increases utility on the other side by more than an agent on the other side would for the original side. Another way of interpreting this result is that since demand and demand elasticities for one side of the market change when the number of agents on the other side increases, the optimal prices charged by a platform may in fact be lower than those predicted if the platform ignored these network effects and assumed demand was fixed. All of these factors lead to more complex pricing dynamics than in traditional one-sided markets, which if ignored can lead to misguided regulatory or antitrust policies (Evans, 2003; Wright, 2004a).

In contrast, if there were no transaction costs or limitations on side-payments, the division of pricing between two sides of a market would not influence transaction volume; instead, any division would merely be a transfer between different sides of the market. In these standard settings, any distinction between how the two sides of the market are charged is unnecessary. Such is the case with a value-added tax, as whether firms pay a value-added tax during the stages of production or consumers pay a sales tax at the point of purchase does not affect transaction volume. Indeed, Rochet and Tirole (2006) use the nonneutrality of price structure—that is, for a given price level, the allocation of prices across sides influences transaction volume—as their definition of what makes a market two-sided.

These insights provided by the literature on two-sided markets naturally apply to the Internet.² First, as discussed, network effects are prevalent: the decision by users to subscribe to Internet access depends on which content providers are online; similarly, the decision of a content provider to invest in the creation of new content and applications depends on the number of users that can be reached. Consequently, the willingness to pay (and elasticity of demand) of an agent varies according to usage by others.

Secondly, there are practical limits on the side-payments that can be easily conducted between different sides of the Internet. Due to informational asymmetries, access or subscription charges to content significantly reduce usage as con-

² Others have explored this connection: for example, Economides and Tåg (2007) offer an application of the two-sided market framework to modeling the Internet, while Hurwitz (2006) and Hemphill (2008) also make the connection in their policy papers.

sumers find it difficult to determine the quality of on-line content before they have used it and thus often choose not to use it at all rather than to pay an up-front fee. Moreover, payments for Internet services in general have a variety of frictions: limited consumer access to credit cards or electronic payment systems; issues with security and fraud; problems of pricing micro-transactions for very limited access; and costs of monitoring whether purchased electronic content is being used only for agreed-upon purposes. Widespread adoption of usage-based pricing by content providers is in many cases thus either impractical or inefficient; it is not surprising, then, that many content providers have chosen to rely on advertising revenues and to provide content at no charge to the user. The presence of these myriad challenges preventing content providers from passing potential termination fees on to consumers thus implies that the division of pricing will be nonneutral. This tempers the “indirect extraction” argument (raised for example by Hemphill, 2008) that a ban on termination fees would be unraveled via side-payments.

Subsidizing Content Creation and Invention

The literature on two-sided markets combined with the economic realities of the Internet suggests that the ban on termination fees serves important economic and potentially social functions.³ First, the rule provides a direct subsidy for the production of content and inventions. Second, it also cheapens market entry by making it easier for users to switch and become content providers themselves. Vinton Cerf (2006), a co-designer of the TCP/IP protocol, captures both these points: “Because the network is neutral, the creators of new Internet content and services need not seek permission from carriers or pay special fees to be seen online. As a result, we have seen an array of unpredictable new offerings . . . [E]ntrepreneurs need not worry about getting permission for their inventions will [sic] reach the end users . . . This is a direct contrast to closed networks like the cable video system, where network owners control what the consumer can see or do.”

Providing subsidies for the production of creative works and innovation is a typical goal of many government programs, including the copyright and patent laws and institutions such as the National Institute of Health or National Endowment of the Arts. Such subsidies are most often justified in economic terms as a remedy for market failure in the production of creative works or new inventions (Landes and Posner, 1989). Since both creative works and inventions have characteristics of a public good, such intervention may be necessary to avoid underproduction.

The pricing structure of the Internet can be seen as an alternative means of subsidizing creativity and innovation (as discussed in Benkler, 1999). As economic analysis suggests, setting a preferable price or ruling out certain types of fees for content providers may encourage creation of content or new inventions that would not otherwise occur. As Lessig and McChesney (2006) note, “more than 60 percent of Web content is created by regular people, not corporations,” and over 100 million

³ Other effects of banning payments from content providers to Internet intermediaries (that is, using a “bill-and-keep” system) can be found in DeGraba (2000), Hemphill (2008), and Wright (2004b).

blogs have so far been documented (*Technorati*, 2008). The Internet, as a platform, has spawned thousands of new firms and millions of sites, from mass content projects such as Wikipedia to search indexers and content aggregators such as Google and Yahoo!.

In addition to the traditional justifications for subsidizing creativity and invention, there are special reasons that subsidizing in this context might be useful. The Internet content and applications market can be understood as a “hit-driven industry,” where hit products like those from Google, Yahoo!, Amazon, eBay, and others create massive spillover value for users. But extreme uncertainty plagues the creation of distinctly new content, and for every hit product there have been numerous failures, with names like Go.com, Pets.com, and Webvan.com. Given that the returns to content production are skewed and the expected value of a new online venture is low, sufficiently low costs of entry may have been and may continue to be crucial.

For similar reasons, the lack of termination fees has also been critical to the “long-tail” model of revenue generation: many Internet businesses rely on an extremely diverse product range that caters to individual niche markets (Anderson, 2006). A diverse collection of websites that yield small value individually but high value when considered as a group might not exist if faced with higher operating costs. Even a more targeted or asymmetric scheme of only levying fees on “successful” or large content providers still has the effect of depressing content creation as it reduces the potential gains to innovation for small entrepreneurs with the dream of making it big.

Of course, for a given price level, subsidizing content comes at the expense of *not* subsidizing users, and subsidizing users could also lead to greater consumer adoption of broadband. It is an open question whether, in subsidizing content, the welfare gains from the invention of the next “killer app” or the addition of new content offset the price reductions consumers might otherwise enjoy or the benefit of expanding service to new users. It may prove useful for economic theory to further illuminate and explore this tradeoff. However, given the possibility of users acting as content providers, a more accurate description of a ban on payments from content providers to Internet intermediaries is this: it is a subsidy to the creative and entrepreneurial at the expense of the passive and consumptive.

The Danger of Fragmentation

Setting aside concerns over incentives for creativity, there is another important concern about allowing Internet service providers to charge fees to content producers: it would almost certainly result in service providers “competing” for content, as seen in other platform industries, by charging different fees and bargaining on exclusive arrangements with content providers. In turn, such bilateral agreements would inevitably lead to fragmentation—where certain content would only be available on certain service providers—and hence multiple “Internets.” For example, cable television is a fragmented network in this sense: not only do users

of cable television face a menu of prices for different numbers of channels, but in addition, certain channels are available only in some geographic areas.

Despite arising in equilibrium, these arrangements need not be efficient in networked industries with externalities and an incomplete contracting space (for example, Segal, 1999; Farrell and Weiser, 2003; Whinston, 2006). Potential welfare losses could also be significant, as consumers would find themselves foreclosed from accessing content available only on rival service providers, and content providers would find themselves unable to reach certain segments of the population captive to service providers with whom no agreement had been reached.⁴

Such arrangements would also be anathema to the principle of universality subscribed to by the designers of the Internet. One visionary of the Internet, Tim Berners-Lee (2008), put it this way: “It is of the utmost importance that, if I connect to the Internet, and you connect to the Internet, that we can then run any Internet application we want, without discrimination as to who we are or what we are doing. We pay for connection to the Net as though it were a cloud which magically delivers our packets. We may pay for a higher or a lower quality of service. We may pay for a service which has the characteristics of being good for video, or quality audio. But we each pay to connect to the Net, but no one can pay for exclusive access to me.” Economists might rephrase this principle of universality in a language of positive network externalities and avoiding fragmentation, but the fundamental social value remains the same.

The Transactions Cost Argument for a Zero-Price

While the two-sided markets analysis might seem to justify setting different prices between agents, a crucial question remains: Why a price of zero?

There is a crucial difference between a low price and a zero-price: the ban on termination fees eliminates an entire class of transaction costs. Similar to the way in which developers can write an application for Microsoft Windows without bargaining with Microsoft at all, content providers can reach all consumers without having to negotiate individually or to pay separate fees to every service provider. For example, assuming Internet access, content providers do not have to negotiate with any service provider to get their initial production started. Transaction costs, of course, can be overcome, but their presence or absence matters.

The most obvious beneficiary of the absence of termination fees and related

⁴ An argument is sometimes made that allowing exclusive arrangements might help new intermediaries to enter network markets (Lee, 2009)—in this case, the argument would be if new Internet service providers could offer exclusive content, it would be easier for them to differentiate themselves from existing providers and gain market share. However, we argue that this argument should not be given great weight in the context of Internet service providers. First, this literature also has noted that a standardization of network platforms often improves welfare (Farrell and Saloner, 1985; David and Greenstein, 1990); in this case, the benefit of standardization arises from maintaining the Internet as the sole platform and ISPs as simply conduits to the greater network. This argument is strengthened given there is no sole ISP, and hence no monopoly rents extracted upon standardization. In addition, just because net neutrality may prevent ISPs from competing on content, they may still compete in other ways: for example, they can differentiate themselves on quality of service.

transaction costs are media forms like blogs; there are millions of content providers for blogs, which are of highly variable quality. If each content provider somehow had to negotiate to gain access to users, the transaction costs alone might endanger their existence in the first place (Heller, 2008). Furthermore, social media sites such as MySpace and Facebook may not have been able to launch before the viability of their sites had been established. For new entrepreneurs or content providers, it has been unnecessary to reach agreements with every carrier to maximize the number of users and contributors accessible—and hence the value of their content—from the start. If content providers had to run a gauntlet of fees before being widely available, many business models would not have been feasible and many content providers may not have entered.

We note the lack of fees for providing Internet content lies in sharp contrast to other networks such as cable television, which involve intensive negotiation over prices for a channel's inclusion in a cable package. In these contexts, such frictions pose significant difficulties for new content in acquiring distribution and obtaining an audience. In a sense, the transaction costs induced by this bargaining often favors the established, well-financed, or overtly commercial at the expense of niche, specialized, and unproven; as we have argued, content exhibiting the latter characteristics may very well have defined much of the Internet's value and worth.

What about Content Providers Charging Service Providers?

We have not yet discussed the possibility of content providers levying fees on Internet service providers. Should these types of reverse charges be allowed? After all, if subsidization is appropriate, why stop at a zero-price and why not have ISPs pay for content? Such a regime would be similar to what is done in the cable industry, where content providers like CNN and ESPN charge cable operators a per-subscriber fee for the right to carry their content. In the Internet context, some content providers have experimented with such fees; for example, ESPN charges service providers for the right to access its ESPN360.com broadband site, as opposed to charging individual users directly (*Wall Street Journal*, 2006).

For reasons similar to those outlined above, we do not think such arrangements are desirable. The use of discriminatory termination fees, even if negative for some content providers, may still lead to positive fees levied on others. But even if asymmetric regulation would be possible—banning one direction of fees—it would still be the case that allowing for any nonzero-pricing would introduce a new class of fees to the Internet and substantially increase transaction costs, favoring some types of content providers over others. Furthermore, such fees also exacerbate the problem of fragmentation and consumer foreclosure: if some Internet service providers did not wish to pay for certain content, it would be difficult to force them to do so.

At the same time, direct regulation for this particular direction of fee payments may not be necessary, for several reasons. First, any content provider that makes a decision to charge fees will necessarily internalize the effect of having fewer consumers that can access its site; as will be discussed later, internalization is not the

case with Internet service providers levying termination fees since a collective action problem is introduced. Second, unless consumers were willing either to switch service providers or to terminate Internet access altogether in response to a particular content provider's availability, a content provider would not be able to extract rents from an ISP; it is difficult to imagine very many (if any) content providers that fit this criteria. Under these conditions, it is unlikely that a significant number of content providers (if any at all) would find it desirable to charge service providers for access. Even ESPN has shown signs of weakening its policy: it now allows anyone with a .edu or .mil domain to access its broadband site, where previously only 20 million subscribers of the ISPs that had signed agreements with ESPN had access (*New York Times*, 2008).

Other Questions and Concerns

Mandating net neutrality via government regulation or even maintaining the current de facto prohibition on termination fees has been controversial. Here, we address a number of questions and concerns that have been raised.

Won't Internet Service Providers Set Appropriate Fees on Their Own?

If setting prices to subsidize content providers produces socially desirable outcomes, wouldn't service providers simply do so? Wouldn't an Internet service provider internalize the externalities across groups and subsidize the "right" side on its own?

One possible answer is "yes" and that this is already happening. Since the early 2000s, despite some early stated interest in charging termination fees to content providers, no Internet service provider has actually tried to do so. Although there have been a few attempts by service providers to limit access to certain types of content, there are relatively few examples. It is unclear, however, whether that behavior is motivated by conscious behavior and internalization of externalities, or by the existing threat of regulation which functions as a form of enforcement.

At the same time, it seems implausible that Internet service providers have appropriate incentives to price according to the social optimum. First, two-sided market theory models predict even a monopolist provider does not subsidize the "right" sides as much as a social planner would, as profit-maximizing prices are higher than those imposed in the social optimum (Armstrong, 2006). Second, and perhaps more interestingly, the fact that a customer will often have a different ISP than a content provider means there will be strong incentives to charge fees even if zero-prices were socially optimal. Furthermore, because the value of content is shared not only by consumers subscribed to a particular ISP but across all consumers, ISPs do not completely internalize the impact of charging termination fees to content providers.

To see this point more clearly, consider the following: assume there are three Internet service providers for users—A, B, and C—and consider A's decision to

charge fees to all content providers wishing to reach its own customers. Service provider A would gain revenue and account for the possibility that such fees would lead to potentially less content (via depressed investment and innovation) for its own users. However, A does not completely internalize how its action would negatively influence content production for the entire Internet, as it does not account for the effect of a reduction in content on users of B and C. Similarly, B and C's incentives to increase termination fees would also be misaligned with the social optimum, since they also fail to completely internalize the effects of their own pricing decisions on users of other ISPs.

In effect, service providers face a prisoner's dilemma: it might be individually optimal for one provider to defect and charge positive fees to content providers, although if all content providers charged such fees, the outcome would be worse than had all providers refrained from doing so. In this sense, the existing *de facto* practice of zero-pricing for content providers on the Internet can be understood as a solution to this collective action problem. Given the temptation to defect, regulation in support of net neutrality—or the threat of such regulation—can play a useful role in maintaining a cooperative solution.⁵

Could Charging Content Providers Help Pay for Network Development or Reduce Web Congestion?

Allowing for termination fees could generate more revenue for service providers; hence, some argue that relaxing the zero-price rule will provide carriers with greater incentives to invest in existing infrastructure. Hemphill (2008) writes of an implicit tradeoff in which “not only content innovation but also infrastructure innovation must be taken into account, and that subsidizing content development necessarily must be at the expense of network development.” If so, perhaps the Internet has now reached a stage of maturity in which, even if innovation in content provision is still desirable, other objectives have become more important?

Edward Whitacre, former CEO of AT&T, made a similar claim (*Business Week*, 2005): “Now what they [content providers] would like to do is use my pipes free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it. So there's going to have to be some mechanism for these people who use these pipes to pay for the portion they're using. Why should they be allowed to use my pipes? The Internet can't be free in that sense, because we and the cable companies have made an investment and for a Google or Yahoo! or Vonage or anybody to expect to use these pipes [for] free is nuts!”

There are two flawed assumptions in this line of argument. First, it is unclear

⁵ Although allowing Internet service providers to form a cooperative may ameliorate the coordination problem, there still is the misalignment of monopolist incentives from the social optimum; furthermore, and perhaps more importantly, the existence of cooperatives may be unpalatable for collusive possibilities that may be encouraged (for example, the credit card industry, which utilizes a cooperative structure, has been the subject of antitrust scrutiny and litigation, including *United States v. Visa U.S.A., Inc. et al.* (344 F.3d 229 [2d Cir. 2003], *cert. denied*, 543 U.S. 811 [2004]).

that the ability to charge termination fees would, in fact, increase incentives to upgrade existing infrastructure. Termination fees may provide a way to increase profits of Internet service providers regardless of whether they upgrade their pipes; the impact on the *marginal* incentive to invest is indeterminate.⁶ Instead of investing in faster or more reliable service, firms could also pay out a (greater) dividend, undertake other projects, or even invest in increasing its returns on existing content by making it scarce and exclusive. As Wu argued (Wu and Yoo, 2007): “If you can generate revenue by charging content providers to reach customers, as opposed to charging for bandwidth, something happens. The incentives become mixed, as the provider gains an incentive to maintain a level of scarcity, and thereby maximize gatekeeper revenue. So I don’t agree . . . that [termination] fees will necessarily spark more last-mile innovation.”⁷

The second problem, which speaks directly to Whitacre’s claim, is that the absence of termination fees does not imply Internet service providers are not properly compensated for the use of their “pipes.” Indeed, ISPs, including AT&T, are actually paid twice. First, any network that transmits a packet through AT&T’s network provides payment to AT&T determined by negotiated peering or transit agreements (and that network was also compensated when it received the original packet from another network, content provider, or end-user). Second, AT&T already charges consumers for access, and nothing in the current net neutrality regime prevents charging higher prices to consumers who utilize more bandwidth or demand faster service. In a sense, claims that content providers should pay for use of a consumer’s ISP when consumers already pay is comparable to the Postal Service demanding a recipient pay to receive a package for which a sender has already paid postage.

Similar reasoning can be used to counter arguments which claim that termination fees might serve to reduce “web congestion.” Content providers, the argument goes, design their applications without taking into account the marginal cost of higher bandwidth usage, and hence “overuse” bandwidth in their designs; that is, these content providers design applications that “spend” bandwidth with wild abandon. However, as discussed, content providers are already forced to take into account the costs of bandwidth usage: those which utilize an Internet service provider must pay the access and usage fees to make the content available in the first place, and those which are connected directly to other network providers pay fees for egress traffic based on existing peering and transit arrangements. The less

⁶ In other words, in a hypothetical market already served by at least one Internet service provider where termination fees are allowed, upgrading pipes would not necessarily lead to an increase in profits. For example, in markets with at least two ISPs, the gains to investment may be competed away. However, termination fees may encourage network expansion to consumers not served by any ISP. In these instances, there may be a tradeoff between expanding Internet access to new markets on the one hand and incentivizing content creation and avoiding Internet fragmentation through a zero-price rule on the other. At the same time, by reducing content creation and availability, termination fees could also lower consumer willingness-to-pay.

⁷ The “last-mile” refers to the technologies and processes (for example, coaxial, wireless, fiber) that connect an ISP to its customers.

bandwidth a content provider uses, the less it pays. Furthermore, Internet service providers can, and in certain circumstances do, charge end-users for the amount of bandwidth used regardless of what content is consumed; thus, content providers that utilize massive amounts of bandwidth will also face and internalize lower demand by consumers.

Would Favoring Some Content Improve Internet Service?

In this section, we consider the possibility that allowing networks to favor some content over others could improve service overall. If hypothetically a network could recognize and prioritize packets more sensitive to delay, like video packets, over packets that are insensitive, like email, the network would in theory function better. Similarly, if packets could be transmitted over shorter distances, less congestion would occur. This reasoning has been the basis for many proposals for “quality of service” enhancements to the Internet since at least the 1990s, and has motivated the rise of content distribution networks and caching services, such as Akamai, which “mirror” content across servers located around the world and thus allow providers to pay for improved delivery of content. Indeed, network management and quality of service inherently requires some form of packet discrimination or content co-location, and are practices with which we do not necessarily take issue (Wu, 2003); in our view, they may be palatable as long as payment is not demanded from content providers by Internet service providers as a requirement for service.⁸

Many global schemes for prioritizing some packets of information over others have so far failed because of a collective action problem inherent in their design. The Internet is comprised of hundreds of Internet service providers and millions of content providers worldwide. So far, agreement on standards to prioritize traffic on the shared network has been impossible to reach, as has any agreement to honor any standards for prioritization. There is an obvious incentive to label every packet as a “high priority” packet on the assumption that everyone else will do the same.

However, it is an open question whether using prices could somehow overcome this problem. One proposal that has been raised to address network congestion in the last mile would be to create a *tiered structure* for consumer ISP traffic: allow *all* content to travel *freely*, but at the same time allow individual Internet service providers to create a “preferred” service for traffic, or a “fast-lane,” for a fee that does not depend on the identity of the content provider. In our view, this approach has the advantage over termination fees by allowing anyone access to faster service—incentivizing content providers to only label traffic “priority” if the cost was deemed worthwhile—while not foreclosing those who still opt for the “free” Internet. As a result, we do not feel as

⁸ We believe caching agreements with Internet service providers and services provided by distribution networks do not impinge on content provision because these services are available to all content providers and content providers that do not use these services are still accessible by consumers. Although these services do improve the performance of certain content providers vis-à-vis others, the services are only worthwhile for content providers with significant traffic and bandwidth demands; for new entrants with low bandwidth requirements, such services provide little benefit and are a nonissue.

though a zero-pricing rule should prohibit this particular implementation, as here content providers are not *forced* to pay a termination fee to access users.

Yet although this particular solution may have desirable benefits, we raise a warning: unless sufficient bandwidth and quality of service can be guaranteed for the “free” Internet, there is a risk that such tiering will serve to sidestep the *de facto* prohibition on termination fees. For example, a priced-priority system could simply become itself a *de facto* fee charged for all content providers if the “free” Internet was of sufficiently poor quality and consumers shifted their usage behavior accordingly. In other words, even if ISPs were prevented from charging discriminatory fees to content providers (thereby reducing bargaining frictions and the potential for fragmentation), tiering still could result in transfers from content providers to Internet service providers. As argued previously, this might dampen the introduction of new content and services and eliminate the subsidy for content innovation currently provided by net neutrality.

Should the Net Neutrality Argument Be Applied to Other Networks?

We have argued the Internet’s history may have created a built-in subsidy for competitive entry and creativity. There is, however, an open question of whether the principle should be extended to other networks. While a full treatment of this question is well beyond the scope of this paper, we do wish to be careful about generalizing our arguments to other industries: in particular, several institutional details may differentiate the Internet from attempts to implement a zero-pricing rule in other networks. For example, one important consequence of the Internet’s universal design is that the bandwidth used by any one content provider is dynamic and proportional to its popularity: only content that is visited or popular consumes common resources, whereas sites that are never accessed utilize zero network bandwidth. Consequently, there is effectively no opportunity cost of subsidizing new content and lowering the barrier to entry, since other content is not precluded from existing or reaching users. In contrast, in media networks such as radio or cable television, each station uses a fixed amount of bandwidth or spectrum regardless of its popularity; similarly, new products sold in stores consume physical space and inventory even if no one purchases them. Thus, even if subsidizing content may be desirable, the scarcity of airtime, spectrum, or shelf space may very well render zero-pricing unappealing and undesirable in other industries.

Concluding Remarks

At its broadest, the net neutrality debate in the United States and around the world is a reincarnation of an age-old debate about the duties of firms that supply infrastructure services essential to the economy, or—in the old common law phrase—firms “affected with the public interest.” In the nineteenth century, trains and canals were the focus of this debate; in the twentieth century, it was the telephone and the electric systems; and in the twenty-first century, the Internet has seized center stage.

This paper has highlighted a potential benefit of the zero-pricing aspect of net neutrality, which prevents Internet service providers from levying termination fees on content providers. The theory of two-sided markets provides an underlying rationalization for how this practice can subsidize the creation of new content and spur innovation while avoiding fragmentation of the Internet. Several open questions remain, including how close the optimal subsidy for content creation is to a zero-price rule, and to what extent welfare gains from increased content production due to a zero-price may be offset by potentially higher access or usage fees charged to consumers.

At the same time, a more fundamental question that underlies this paper is what, if anything, sets the Internet apart from other networks, past and present? This question suggests a much broader agenda for research: namely, understanding in a more parsimonious manner how different pricing rules and other features of information networks affect their influence. We have mentioned two differences—a *de facto* ban on termination fees, and a rough proportionality between content popularity and bandwidth usage—that set the Internet apart from the other networks of our time, including telephone, cable TV, and broadcasting networks. Yet while this much may be clear, we do not have anything close to a full vocabulary for understanding the different choices implicit in the designs of different networks. And we have an even weaker understanding of what the larger effects of such choices will be. Although in this paper we have isolated one interesting effect—namely, than a ban on termination fees can be used to encourage market entry by creators and innovators—this point is far from a full understanding of networks and their larger effects on society and the world.

■ *We thank the editors and Scott Hemphill for their helpful comments.*

References

- Anderson, Chris.** 2006. *The Long Tail*. New York: Hyperion.
- Armstrong, Mark.** 2006. "Competition in Two-Sided Markets." *RAND Journal of Economics*, 37(3): 668–91.
- Benkler, Yochai.** 1999. "Free as the Air to Common Use: First Amendment Constraints on Enclosure of the Public Domain." *New York University Law Review*, 74(2): 354–446.
- Berners-Lee, Tim.** 2008. "Neutrality of the Net." <http://dig.csail.mit.edu/breadcrumbs/node/132> (accessed on November 11, 2008).
- Business Week.** 2005. "Online Extra: At SBC, It's All About 'Scale and Scope.'" November 7. http://www.businessweek.com/magazine/content/05_45/b3958092.htm.
- ▶ **Caillaud, Bernard, and Bruno Jullien.** 2003. "Chicken & Egg: Competition among Intermediation Service Providers." *RAND Journal of Economics*, 34(2): 309–28.
- Cerf, Vinton G.** 2006. Prepared Statement, U.S. Senate Committee on Commerce, Science, and Transportation Hearing on "Network Neutrality." <http://commerce.senate.gov/pdf/cerf-020706.pdf>.
- ▶ **David, Paul, and Shane Greenstein.** 1990.

"The Economics of Compatibility Standards: An Introduction to Recent Research." *Economics of Innovation and New Technology*, 1(1–2): 3–41.

DeGraba, Patrick. 2000. "Bill and Keep at the Central Office as the Efficient Interconnection Regime." OPP Working Paper No. 33, Office of Plans and Policy, Federal Communications Commission, Washington DC.

Economides, Nicholas, and Joacim Tåg. 2007. "Net Neutrality on the Internet: A Two-sided Market Analysis." NET Institute Working Paper No. 07-45.

Evans, David S. 2003. "Some Empirical Aspects of Multi-sided Platform Industries." *Review of Network Economics*, 2(3):191–209.

▶ **Farrell, Joseph, and Garth Saloner.** 1985. "Standardization, Compatibility, and Innovation." *The RAND Journal of Economics*, 16(1): 70–83.

Farrell, Joseph, and Philip J. Weiser. 2003. "Modularity, Vertical Integration, and Open Access Policies: Towards a Convergence of Antitrust and Regulation in the Internet Age." *Harvard Journal of Law & Technology*, 17(1): 85–134.

Hagiu, Andrei. 2006. "Pricing and Commitment by Two-Sided Platforms." *RAND Journal of Economics*, 37(3): 720–37.

Heller, Michael A. 2008. *The Gridlock Economy*. New York: Basic Books.

Hemphill, C. Scott. 2008. "Network Neutrality and the False Promise of Zero-Price Regulation." *Yale Journal on Regulation*, 25(2): 135–180.

Hurwitz, Justin. 2006. "Neighbor Billing and Network Neutrality." *Virginia Journal of Law and Technology*, 11(9): 1–33.

▶ **Katz, Michael, and Carl Shapiro.** 1986. "Technology Adoption in the Presence of Network Externalities." *Journal of Political Economy*, 94(4): 822–41.

▶ **Landes, William M., and Richard A. Posner.** 1989. "An Economic Analysis of Copyright Law." *Journal of Legal Studies*, 18(2): 325–63.

Lee, Robin S. 2009. "Vertical Integration and Exclusivity in Platform and Two-Sided Markets." <http://pages.stern.nyu.edu/~rslee/papers/VIExclusivity.pdf>.

Lessig, Lawrence, and Robert W. McChesney. 2006. "No Tolls on the Internet." *Washington Post*, June 8, A23.

Licklider, Joseph C. R., and Robert W. Taylor. 1968. "The Computer as a Communications Device." *Science and Technology: For the Technical Men in Management*, April, no. 76, pp. 21–31. (Re-

printed in 1990 in *In Memoriam: J. C. R. Licklider: 1915–1990*, Report 61, pp. 21–41. Palo Alto, CA: Systems Research Center, Digital Equipment Corporation.)

Liebowitz, S. J., and Stephen E. Margolis. 1994. "Network Externalities: An Uncommon Tragedy." *Journal of Economic Perspectives*, 8(2): 133–50.

The New York Times. 2008. "ESPN to Offer Sports Events on the Web Free to Some." February 4. <http://www.nytimes.com/2008/02/04/business/media/04espn.html> (accessed November 10, 2008).

Posner, Richard. 1998. *Economic Analysis of Law*. 5th ed., pp. 36–50. New York: Aspen.

Reardon, Marguerite. 2006. "FCC approves AT&T–BellSouth merger." CNET News, December 29. http://news.cnet.com/FCC-approves-ATT-BellSouth-merger/2100-1036_3-6146369.html.

▶ **Rochet, Jean-Charles, and Jean Tirole.** 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association*, 1(4): 990–1029.

Rochet, Jean-Charles, and Jean Tirole. 2006. "Two-Sided Markets: A Progress Report." *RAND Journal of Economics*, 37(3): 645–67.

▶ **Segal, Ilya.** 1999. "Contracting with Externalities." *Quarterly Journal of Economics*, 114(2): 337–88.

Technorati. 2008. "State of the Blogosphere." <http://technorati.com/blogging/state-of-the-blogosphere> (accessed on November 12, 2008).

The Wall Street Journal. 2006. "ESPN Charges Net Providers For Right to Offer Broadband Web Site." August 1. <http://online.wsj.com/article/SB115439535367922979.html>.

Whinston, Michael D. 2006. *Lectures on Antitrust Economics*. Cambridge: MIT Press.

Wright, Julian. 2004a. "One-sided Logic in Two-sided Markets." *Review of Network Economics*, 3(1): 44–64.

▶ **Wright, Julian.** 2004b. "Pricing Access to Internet Service Providers." *Information Economics and Policy*, 16(3): 459–73.

Wu, Tim. 2003. "Network Neutrality, Broadband Discrimination." *Journal of Telecommunications and High Technology Law*, vol. 2, pp. 141–79.

Wu, Tim, and Christopher Yoo. 2007. "Keeping the Internet Neutral?: Tim Wu and Christopher Yoo Debate." *Federal Communications Law Journal*, 59(3): 575–92.