NETWORK NEUTRALITY: REGULATING WITH MORE QUESTIONS THAN ANSWERS

HOWARD A. SHELANSKI*

INTRODUCTION

"Network neutrality," while subject to varying definitions, can be summed up as the principle that "all like Internet content must be treated alike and move at the same speed over the network. The owners of the Internet's wires cannot discriminate."[1] The policy implication is that network operators should not be allowed to "create different tiers of online service" by selling different levels of access at different prices to different providers of on-line content and services.[2]

Proposals for network neutrality regulation have sparked particularly intense debate. Advocates and opponents of regulation have each predicted dire consequences from, respectively, leaving networks free to vary terms of access they offer to upstream providers of content and services[3] or restricting them from doing so. As the debate has continued between those who argue that network neutrality regulation is necessary to preserve applications innovation and those who argue that such regulation would harm the growth and development of underlying network infrastructure, Congress has been awash with legislative proposals from both perspectives.[4]

---

* Professor of Law, University of California at Berkeley. This essay is based on the author's presentation at the Silicon Flatirons Digital Broadband Migration Conference, February 19-20, 2006, University of Colorado, Boulder. The author is grateful to Joe Farrell, Larry Lessig, Jim Speta, Barbara van Schewick, and Phil Weiser for helpful comments and discussions.

1. Lawrence Lessig & Robert W. McChesney, *No Tolls on the Internet*, WASH. POST, June 8, 2006, at A23, *available at* http://www.washingtonpost.com/wp-dyn/content/article/2006/06/07/AR2006060702108.html.

2. *Id.*

3. I will refer to providers of Internet content and services generically as "applications providers" for the rest of this essay.

4. *See* ROBERT D. ATKINSON & PHILIP J. WEISER, INFO. TECH. & INNOVATION FOUND., A "THIRD WAY" ON NETWORK NEUTRALITY 2 n.3 (2006), http://www.itif.org/files/netneutrality.pdf.

Why such attention to network neutrality? The reason may lie in the fact that, although vertical issues have long been central to telephone regulation,[5] the stakes for consumers have changed with the Internet. Only a few years ago, the principal value of the telephone network to consumers was person-to-person voice communication and the principal value of cable networks was video programming. Complementary, vertical services like voice mail or information services were comparably modest in importance. Now, those same networks deliver a vast universe of content and services through the Internet. Some such services, for example Internet telephony ("VoIP") or video services ("IP-TV"), may compete directly with the core services of the underlying networks. But most services are complements, not competitors, to the networks over which consumers reach the Internet, and there is enormous value in those complementary applications. Telephone and cable networks have gone from wagging the tail to wagging the dog with respect to vertical services and their importance to consumers. While the increasing value of the applications market gives rise to concern over vertical discrimination, it simultaneously raises the potential benefits of vertical relationships between networks and applications providers.[6] Particularly for new and commercially risky applications, vertical relationships can, at least theoretically, reduce transaction costs and bring new products and services to market faster. Not surprisingly, therefore, network neutrality regulation has both its advocates and opponents who speak in adamant terms about the consequences of either allowing network owners to discriminate among applications providers or barring them from doing so.

Proponents of regulation confidently argue that discriminatory access terms will chill innovation at the edge of the network, reducing the flow of new services and applications for consumers.[7] Opponents argue with equal force that a ban on discrimination will dampen innovation and investment in the core of the network, reducing capacity and shifting costs to consumers.[8] Applications providers argue that

---

5. *See* STUART MINOR BENJAMIN ET AL., TELECOMMUNICATIONS LAW AND POLICY chs. 13-14 (2d ed. 2006).

6. *See, e.g.*, Oliver E. Williamson, *Assessing Vertical Market Restrictions: Antitrust Ramifications of the Transaction Cost Approach*, 127 U. PA. L. REV. 953 (1979); Oliver E. Williamson, *The Vertical Integration of Production: Market Failure Considerations*, AM. ECON. REV., May 1971, at 112.

7. *See, e.g.*, *Net Neutrality: Hearing Before the S. Comm. on Commerce, Science, and Transportation*, 109th Cong. 54-59 (2006) (prepared statement of Lawrence Lessig), *available at* http://www.lessig.org/blog/archives/lessig_testimony_2.pdf.

8. *See, e.g.*, Christopher Yoo, *Beyond Network Neutrality*, 19 HARV. J.L. & TECH. 1, 27-28 (2005); Bruce M. Owen, *The Network Neutrality Debate: 25 Years after AT&T v. United States and 120 Years After the Act to Regulate Commerce*, PERSP. FROM FSF SCHOLARS (Free State Found., Potomac, Md.), Feb. 20, 2007,

discriminatory pricing will unfairly target deep-pocket providers or firms that compete with the platform's own vertical services.[9] Platform providers argue that they have no incentive to make the Internet less attractive to their subscribers and that successful applications providers are free riding on their networks.[10] Each side claims to champion competition and innovation while portraying the other as being something between an opportunist and a gangster.[11] Upon closer inspection, however, each side's arguments beg important questions to which answers are both empirically and theoretically elusive. Those open questions, in turn, weaken the basis for either the outright ban on discrimination sought by network neutrality advocates or the pure laissez-faire sought by its opponents.

This essay will briefly examine several unanswered questions central to the network neutrality debate and discuss their implications for broadband policy. Part I of this article will examine the main claims made by each side of the network neutrality debate and discuss the unanswered questions upon which the merits of those arguments depend. Part II will analyze the policy implications of those unanswered questions, examine the balance of risks at issue in network neutrality regulation, and discuss how policy should account for those risks in the presence of incomplete information.

I. STRONG ASSUMPTIONS ABOUT REGULATING (OR NOT)

Proponents of network neutrality regulation contend that discriminatory network access terms will selectively impede applications providers' access to consumers and thereby chill innovation at the edge of the network (meaning innovation by those who use the network as a medium for providing their content and services to consumers), reducing the flow of new services and applications to the market. They contend that discrimination would force potential innovators either to buy a costly level of access or risk providing a second-class service with reduced priority to the conduits that reach consumers and, in turn, reduced chances for commercial success. Either choice imposes costs that will cause applications developers on the margin to engage in less innovation.

---

http://www.freestatefoundation.org/images/The_Net_Neutrality_Debate-Bruce_Owen.pdf.

9. *See* Letter from Jeff Bezos, Founder & CEO, Amazon.com, et al., to Joe Barton, Chairman of U.S. H. Comm. on Energy and Commerce et al. (Apr. 5, 2006), *available at* http://markey.house.gov/docs/telecomm/CEO%20Letter.pdf.

10. Online Extra, *At SBC, It's All About "Scale and Scope"*, BUS. WK., Nov. 7, 2005 (quoting SBC CEO Edward Whitacre on free riding by applications providers), *available at* http://www.businessweek.com/@@n34h*IUQu7KtOwgA/magazine/content/05_45/b3958092.htm.

11. Tim Wu, *Why You Should Care About Network Neutrality: The Future of the Internet Depends on It!*, SLATE, May 1, 2006, http://www.slate.com/id/2140850/.

Advocates thus argue that a level, or neutral, playing field for all applications providers is necessary to preserve the ability of intelligence at the "edge" of the network to drive innovation and increase the welfare of consumers.

Arguments against network neutrality often rest on the similar, but diametrically opposed, proposition that investment and innovation will suffer unless network owners can recover costs imposed by high-volume applications. The innovation at issue here is not at the edge of the network but at its "core." At issue is the need for capacity, reliability, and security for traffic moving across the network. Some network owners argue that the content and service providers whose applications generate the traffic should pay for the capacity to carry it to end users. From this perspective, applications providers impose costs on networks and should bear them accordingly, not shift them to network owners or subscribers. Network operators argue that they have no incentive or ability to exclude or reduce the appeal to consumers of any upstream applications, because those applications are what attract subscribers to their networks. They also note that some applications innovators on the edge of the network might be deterred not by discrimination, but by neutrality, because they will be unable to secure priority access from the network operator for services that need to run with a particular assured quality.

Each set of arguments above raises difficult empirical and theoretical questions, and each depends to some extent on the competitive dynamics of the network access market. The more networks there are in competition with each other for subscribers, the less easily can any individual network engage in inefficient discrimination against particular applications or applications providers. Consumers will choose networks that get them the content and services they want fast and reliably. Which side of the debate one credits will therefore depend, at least in part, on one's view of how competitive the market is and will be.

A. Discriminatory Access and Applications Innovation

Even assuming all applications innovation to be welfare improving, what basis is there for determining how much, if any, innovation deterrence would result from discrimination by platforms in terms of access offered to applications providers? Two proponents of network neutrality regulation offer the following empirical motivation for their claim that non-neutrality would deter innovation:

> More than 60 percent of Web content is created by regular people, not corporations. . . . Most of the great innovators in the history of the Internet started out in their garages with great ideas and little capital. This is no accident. Network neutrality protections minimized

> control by the network owners, maximized competition and invited outsiders in to innovate. Net neutrality guaranteed a free and competitive market for Internet content. The benefits are extraordinary and undeniable.[12]

Taking the above argument to be true, the fact that innovators thrived under a neutral regime does not itself tell us how many of those innovators would have been deterred had network operators offered a tiered set of offerings in which quality rose with price. The empirical observation that has motivated some to advocate network neutrality thus does not necessarily supply empirical support for the innovation deterrence argument on which that advocacy largely rests.

Nor is the logical or theoretical connection between neutrality and applications innovation so clear that the network neutrality advocates' innovation-deterrence argument should be accepted as a matter of reason. First, at least some applications providers may be deterred by the absence of a high-priority tier of access. Some services, for example video services, may need reduced latency to work well, and absence of an assured level of priority raises the risk that such services will fail to live up to their billing, hence deterring their introduction.

Second, there is no reason to assume that most services will in fact be harmed if they are transmitted with the base (*i.e.* lower) level of priority. Comparatively low-bandwidth applications may work perfectly well at lower tiers of access and their innovation might not depend on neutrality. Moreover, even if there is some quality effect, consumers have shown a willingness to tolerate slower interactions on the Internet in return for lower subscription prices. Success of an application, therefore, may not depend on purchasing a costlier tier of access from network operators, especially where there is some way to compensate consumers for any delays in service.

Third, even if neutrality was a causal factor in the explosion of innovation from the edge of the network in the first decade of the commercial Internet, that same environment need not be optimal for the next decade of a more mature Internet. It bears noting that in key areas of commerce, content, and applications, the on-line world is populated by a handful of major players. The brand-name recognition, installed base of customers, and network externalities accumulated by established on-line players could present much greater obstacles in some lines of internet applications than would discriminatory access terms. Indeed, it is precisely the established players who fear non-neutrality because they may be natural, deep-pocket targets for aggressive access negotiation by network operators. Neutrality regulations would protect them from such

12. Lessig & McChesney, *supra* note 1.

pressure.

Neutrality may, however, also benefit established players in another way, this one less sympathetic or potentially beneficial for innovation: access quality may be an important way for new competition in some services to differentiate themselves from incumbents. Established applications providers have little interest in defending against entrants on new competitive dimensions. The "neutral" status quo may therefore be of competitive advantage to applications incumbents while denying a competitive tool to new innovators from the edge.

Finally, platform competition was less developed during the early years of the commercial Internet. Few Americans (19 percent) even had Internet access at all from their homes in 1996, while today most have computers and a choice of broadband access providers.[13] Even if neutrality was necessary to speed applications innovation under the early years of limited broadband availability and no choice of broadband providers, it is unclear that it would be in today's more competitive environment.

The arguments made above do not refute the possibility that non-neutrality will deter applications innovation. They do, however, show that there is little reason to presume such an effect for policy purposes and good reason to question whether non-neutrality will cause the severe harms that some network neutrality proponents suggest. The case for such harmful effects diminishes with increased network competition. Under duopoly, the case is ambiguous. As wireless platforms enter the market to compete against the cable and telephone networks, the ability of any network to discriminate inefficiently by artificially slowing selected traffic to sell priority declines because its rivals will have incentives to offer consumers greater assurance of fast content delivery.

B. Networks and Incentives to Discriminate

Consider next the incentives of network owners to engage in discrimination that harms innovation or consumer welfare. Opponents of network neutrality regulation have argued that network owners would have no incentive to discriminate against applications providers in a way that made network subscription less attractive to consumers. Underlying this claim is the idea that "a monopolist—which, by definition, would have the ability to impede competition in adjacent markets—generally will have no incentive to do so" because it cannot enlarge its profits by doing so.[14] Any reduction in value (or increase in price) of the upstream

13. Press Release, FCC, Federal Communications Commission Releases Study on Telephone Trends (June 21, 2005), *available at* http://www.fcc.gov/Bureaus/Common_Carrier/Reports/FCC-State_Link/IAD/trend605.pdf.

14. James B. Speta, *The Vertical Dimension of Cable Open Access*, 71 U. COLO. L. REV.

application will be met by a corresponding reduction in demand (or decrease in profits) for platform subscriptions, a phenomenon that Joseph Farrell and Philip Weiser have labeled "*internalizing complementary efficiencies*" or "ICE."[15]

Farrell and Weiser demonstrate, however, that while ICE often holds, under many conditions it does not. As Farrell explains, platform owners can often raise their profits by price discrimination, and even if one assumes the price discrimination itself to be efficient (which is not always the case), platform owners may discriminate against providers of complementary services in order to facilitate price discrimination.[16] Farrell illustrates his point through the simple example of a copy machine manufacturer that wishes to price discriminate by selling the copier at a low price and metering use through sale of repair services.[17] In order for repair services to be a metering mechanism for price discrimination, the copier manufacturer must receive revenues for all repairs done to its copiers. One way the manufacturer can do this is to withhold spare parts from independent repair firms and to do all the repairs itself, eliminating competition and reducing efficiency in the complementary repair market. Thus, the non-neutrality of the mechanism used to accomplish price discrimination can involve what Farrell has termed "collateral-damage inefficiency."[18] The important point is that whether or not the underlying price discrimination is itself efficient, that discrimination can be profitable for the manufacturer despite any collateral-damage inefficiency it might cause.

In theory, the manufacturer could avoid this collateral damage through other means of metering. For example, instead of making repair services the metric, the manufacturer could make spare parts the metric and then meter usage of the copier through sales of spare parts to all providers of repair services. Copier owners would retain their choice of service providers and the most efficient service providers would remain able to compete for repair business. To the extent that more efficient metering mechanisms are harder to administer than preemption of competition in the complementary market, however, firms may opt for the latter despite the inefficiency.[19]

975, 997 (2000).

15. Joseph Farrell & Philip J. Weiser, *Modularity, Vertical Integration, and Open Access Policies: Towards a Convergence of Antitrust and Regulation in the Internet Age*, 17 HARV. J.L. & TECH. 85, 89 (2003).

16. Joseph Farrell, *Open Access Arguments: Why Confidence is Misplaced*, *in* NET NEUTRALITY OR NET NEUTERING: SHOULD BROADBAND INTERNET SERVICES BE REGULATED? 195, 199 (Thomas M. Lenard & Randolph J. May eds., 2006).

17. Farrell does not use this example in his 2006 paper but did so in discussions with the author.

18. Farrell, *supra* note 16, at 199.

19. *See, e.g.*, Eastman Kodak Co. v. Image Technical Servs., Inc., 504 U.S. 451, 478

In the context of network neutrality, the pursuit of price discrimination could lead to harmful departures from neutrality toward upstream applications. While Farrell and Weiser show that platform owners may have incentives to discriminate inefficiently where the application competes with a core service of the platform[20] (*e.g.* VoIP for telephone networks or video-on-demand for cable networks), harm may still result even when the upstream application is not one that rivals the platform's main line of business. For example, one mechanism a cable network owner could use to price discriminate is to bundle Internet access with some application, say IP telephony. The network could offer consumers two choices: Internet access for $30 per month, or Internet access for $25 per month if the consumer also subscribes to the network operator for IP telephone service. To make this bundle profitable, the network operator not bound by network neutrality rules might discriminate in the terms of access it provides to rival IP telephone providers to put them at a competitive disadvantage. So long as the increased attractiveness of Internet subscriptions due to the $5 discount outweighs the decrease in attractiveness due to the reduced choice of IP telephone services, the network operator may find the collateral damage to the upstream applications market nonetheless to be profitable. The same scenario could hold for other means of price discrimination, say a phone company's metering of subscribers' Internet usage through video downloads or some other application susceptible to incremental charges.

This is not to say that there are no possible welfare benefits from the price discrimination described above. By using discrimination — whether through bundling, metering, or some other mechanism — to extract high surplus from one set of users, a network operator may enable another set of users to have access where they would not under a single-price regime. This is particularly so in the case for high-fixed-cost services like Internet access, where price discrimination might allow a network to offer some subscribers access at prices closer to marginal cost because it is recovering its fixed costs from other, higher-paying, customers. It is this very ambiguity in the welfare effects of price discrimination and in the incentives to discriminate inefficiently that is important. The welfare ambiguity means that any rule patently barring discrimination could have unintended, negative consequences because the conduct sought to be barred — price discrimination — is neither always bad nor always good.

(1992) (alleging vertical foreclosure by Kodak as a means to leverage profits).

20. *See* Farrell & Weiser, *supra* note 15, at 108.

C. Capacity, Efficient Priority Choices, and Network Investment

A third set of questions in the network neutrality debate revolves around network capacity. If capacity is not scarce, then there is no need for networks to prioritize one provider's traffic over another and no need for investment in new capacity. Capacity thus implicates two important issues for the network neutrality debate. The first is whether upstream price discrimination is necessary to establish priority; the second is whether upstream price discrimination is necessary to recover the costs of investing in new capacity and network technology. The threshold question underlying both of these questions is whether capacity is scarce such that congestion will at least sometimes occur and require networks to prioritize one packet of information over another. If not, then it is hard to see what good could emerge from departures from neutrality, as such departures could be aimed neither at efficiently prioritizing traffic, nor at efficiently recovering network investment.

There may be little agreement over the exact extent of current or future capacity constraints on broadband networks, but neither is there evidence that capacity is so plentiful that congestion, and hence the issue of priority, never arises. Indeed, one report argues that new capacity investment is necessary and that the market does not currently provide adequate incentives for network owners to make such investments.[21] The head of television technology for one of the strongest advocates of network neutrality, Google, in a widely reported statement also emphasized the need for core investment when he said "[t]he Web infrastructure and even Google's (infrastructure) doesn't scale. It's not going to offer the quality of service that consumers expect."[22] Given that capacity constraints cannot be assumed away in the network neutrality debate,[23] the question becomes whether they can supply any justification for differentiating among applications providers in the terms of network access.

One rationale for allowing price discrimination is that it provides a basis for deciding which packet should take priority over another. This is exactly what raises concern among network neutrality advocates; new

---

21. DELOITTE TOUCHE TOHMATSU, TELECOMMUNICATIONS PREDICTIONS: TMT TRENDS 2007 8 (2007), *available at* http://www.deloitte.com/dtt/cda/doc/content/us_tmt_%202007_Telecom_Predictions_011606.pdf.

22. *Google and Cable Firms Warn of Risks From Web TV*, USA TODAY, Feb. 7, 2007 (quoting Vincent Dureau), *available at* http://www.usatoday.com/tech/news/2007-02-07-google-web-tv_x.htm.

23. Indeed, such an assumption implies either that the marginal value of investment in the core platform infrastructure is zero or that it is always lower than the marginal value of applications innovation. As discussed below, there are many unknowns about the relevant incentives to innovate and about the marginal benefits to consumers of different innovations; but the evidence suggests that the core cannot simply be ignored in favor of the edge.

applications providers will have to either pay or sit in line.[24] As discussed above, charging for priority may or may not have a significant negative impact on applications innovation. But if there really is a need to prioritize, it is important to examine the alternatives before ruling out price mechanisms. The most neutral alternative of random selection would serve consumers poorly. A spam e-mail is likely to be less valuable to either consumer or provider than a VoIP call or a paid music download. Random selection could lead the spam to be delivered first, however, benefiting no one except the provider of the lower-value service.

A more nuanced alternative is suggested by the definition of network neutrality at the beginning of this article: "all *like* . . . content must be treated alike and move at the same speed."[25] Under a close reading of this definition, it might be fine for the network to prioritize VoIP over e-mail, so long as all VoIP were treated the same and all e-mail were treated the same. While such hierarchy of uses might be better than random prioritization, it still raises potential problems because it puts the network owner in the position of having to decide which uses or categories of content should be prioritized over others, which uses are "like" other uses, and where innovative new uses should be placed in the priority queue. Defining a clear and administrable regulatory standard for "like content" will prove difficult.

Creating a market for priority can alleviate the difficulties with random or "like use" prioritization and reduce the allocative inefficiency that can result from those mechanisms. Network investment could become more efficient because firms with a desire for priority will capture direct private benefits (less delay for their packets) of their payments to the network operator. When the network owner or subscribers must bear the costs, the benefits are more diffuse, creating the potential for underinvestment. Moreover, to the extent price discrimination allows more highly valued information to move faster, it has the potential to increase the efficiency and consumer welfare of Internet activity. On the other hand, to the extent price discrimination is used in a targeted way as an anticompetitive strategy to raise the costs of particular applications providers, it can be harmful. Again, the non-neutral strategy can have either (or both) positive and negative effects.

The next question related to capacity is whether recovery of capacity investment supplies a rationale for price discrimination toward applications providers. Networks receive revenues from subscribers, raising the question of why they would need to charge applications providers for access. There are several reasons why recovering network

24. *See*, *e.g.*, Wu, *supra* note 11, at 3.
25. Lessig & McChesney, *supra* note 1 (emphasis added).

costs from subscribers alone might not be optimal. First, even though networks can and do charge subscribers different monthly fees for different Internet access speeds, that pricing mechanism may leave some subscribers who are willing to pay the cost of higher-speed access nonetheless unwilling to pay its price. Within each tier of access, there will be relatively high-usage subscribers and relatively low-usage subscribers. Because all subscribers to a given tier pay the same price, the latter may pay for more speed and capacity then they use while the former pay for less than they use. The subscription price that the relatively low-usage consumers pay is therefore above the costs they impose on the network. Were the subscription price for these users lower and more reflective of their actual usage, they would attract yet lower-usage customers whose willingness to pay was above cost, but not quite up to the existing monthly charge for the higher tier of access. To the extent payments from applications providers can ameliorate this potential inefficiency of consumer-side charges, charging those applications providers can be beneficial.

Second, even if subscription rates can be structured better to reflect each subscriber's actual usage, there may still be inefficiency in on-line consumption. One reason stems from the costs of trying out new, high-bandwidth content and applications. If consumers are paying the full costs of their usage, they may hesitate to try new services that would increase their costs. Some kind of transfer payment from the applications providers to consumers could overcome this inefficiency, although such compensation mechanisms might involve high transaction costs. If applications providers would be willing to pay more to networks in return for subscribers who have faster connections and are more willing to consume various content and services, then it might be more efficient, as well as more profitable, for networks to reduce subscription prices in conjunction with charging applications providers for different levels of access.

Finally, consumers and applications providers may have asymmetric valuations of their interactions. It may be more valuable for applications providers to have consumers use their services than it is for consumers to receive them. This is particularly true where the applications provider is paid by a third party — perhaps an advertiser or search listing — based on the number of people who visit the site. Any given consumer might find the experience worthless and merely "click through" the site. The applications provider may, however, benefit from that very same click-through and therefore, have an interest in reducing the cost to subscribers of accessing their sites. If the network can only charge the consumer for network access, the joint surplus of consumers and applications providers might be lower than it would be if applications providers could pay to speed interactions with, and perhaps reduce prices to, consumers.

The above three reasons why it might not be efficient to charge only subscribers for use of network infrastructure do not resolve the question of whether price discrimination toward applications providers will improve consumer welfare or efficiency. They do show, however, that this issue is complex and that arguments for upstream price discrimination cannot be ignored just because networks already charge subscribers. Internet platforms may well have the attributes of two-sided markets, in which charging end-users and applications providers can be more efficient than placing the charges on one side alone.[26] Whether or not they do, and whether or not the gains from two-sided pricing offset possible costs, are beyond the scope of this paper and are important topics for further research. For current purposes, however, the important point is that the question of the comparative costs and benefits of one-sided versus two-sided pricing is an open one that should not be assumed away on either side of the network neutrality debate.

II. COMPARATIVE RISKS OF ALTERNATIVE FORMS OF NON-NEUTRALITY

The previous section demonstrates that the effects of network non-neutrality toward applications providers are ambiguous, with some possibility that neutrality could deter applications innovation but some possibility too that it could benefit, to varying degrees, network investment, applications competition, and allocative efficiency. Conversely, mandatory neutrality could benefit applications innovation and prevent collateral inefficiencies due to anticompetitive vertical discrimination, but could also reduce the efficiency of investment and the volume and nature of on-line transactions. In neither case, however, are the benefits either sufficiently sure or substantial to justify a policy that pursues one set of objectives (*e.g.* applications innovation) to the exclusion of others (*e.g.* network investment). There are too many open questions about the impact of either laissez-faire or a strict neutrality rule to make a persuasive case for either solution. Either choice is uncertain to achieve its intended objectives and likely to involve tradeoffs and to entail a balance of risks with respect to other beneficial objectives.

This section argues that the policy choice need not be as stark as that between complete neutrality and unrestrained laissez-faire. Discrimination varies in its motivations and methods, and different kinds of network discrimination differ in the balance of risks they entail for networks, applications providers, and consumers. Regulation that restricts some forms of discrimination but not others might protect against the worst harms of non-neutrality without eliminating some of

26. *See* Jean-Charles Rochet & Jean Tirole, *Two-Sided Markets: A Progress Report*, 37 RAND J. ECON. 645 (2006), *available at* http://idei.fr/doc/wp/2005/2sided_markets.pdf.

the investment and efficiency benefits that differentiated access terms for applications providers might allow.

A. Reasons for a Network to Discriminate

Several things might motivate a firm to discriminate in the terms it offers to customers or providers of complements. At the broadest level, a firm might discriminate because it must due to scarcity. In the network context, a firm might be driven to sell priority because congestion requires packets to be dropped at times. In such a case, discrimination could take the weak form of granting priority to some packets only when the capacity constraint binds. An analogy might be a traffic lane that is reserved for eligible vehicles only at rush hour, but is open to general use at other times. This kind of discrimination is what Edward Felten calls "minimal discrimination."[27]

Alternatively, a firm might sell priority because it can manipulate traffic in either beneficial or harmful ways. The analogy here is to a special traffic lane that is reserved all the time, even at times when there would be no congestion were that lane open to use by all. The result could be to raise the probability of delay on the non-reserved lanes, thus attracting customers who won't risk moving slowly and want an assurance of moving quickly at all times. This kind of discrimination is what Felten calls "non-minimal" or "delay" discrimination.[28] Such discretionary prioritization is not necessarily inefficient, depending on the relative costs of delay to those users that incur the delay and those that pay to avoid it. It does, however, raise the prospect of inefficiency and anticompetitive manipulation. Even at this general level there are different risks of harm to competition and innovation. Discrimination driven by necessity that occurs only when capacity constraints bind runs less of a risk of harm than discrimination that is driven by market power and the ability to manipulate traffic.

Discrimination could be further motivated by a number of more specific forces that work in tandem with those motivations discussed above. For example, a network could discriminate against an applications provider as part of an anticompetitive strategy to harm an application or provider that the network does not like, perhaps to shift market share of a complement to the network operator. Alternatively, the network could discriminate because it realizes that some providers are willing to pay more if pushed to do so, thus shifting surplus from the applications provider to the network operator.[29] Or, the network could

27. Edward Felten, *Nuts and Bolts of Network Neutrality*, 6 J. ON TELECOMM. & HIGH TECH. L. (forthcoming 2008).

28. *Id.*

29. Such arguments are sometimes framed as a claim that some applications providers are

price discriminate to recover operating expenses or investment from those applications providers who cause the network to incur higher costs, thus shifting costs from the network operator to the applications provider. Again, each of these motivations entails different risks to competition and innovation, with raising rivals' costs being the most harmful motivation and cost-recovery being the most consonant with competition and innovation.

B. Methods of Network Discrimination

Next, consider alternative methods of discrimination. An important distinction is between targeted and non-targeted price discrimination. In broad terms, a network operator could select particular users or uses that it thinks should pay more for access and adopt policies that induce those firms to do so. For example, a network operator could set a higher price for all streaming video providers on the ground that such providers use a lot of platform capacity. The network could give other uses priority over the packets of any streaming video provider that fails to buy the higher level of access. Alternatively, the network could simply sell priority to whomever wants it, leaving each streaming video provider (or provider of any kind of any application) to decide for itself whether it is willing to have its packets delayed when there is congestion. The competitive risks vary for different kinds of targeted and non-targeted pricing.

Targeted and non-targeted pricing can also take several forms. A network operator could differentiate in its access pricing among specific users, particular kinds of use, or amounts of usage. The first, the targeting of specific users, would set prices depending on the identity of the provider whose traffic is moving over the network. Such categorization could simply be a proxy for use or usage. For example, if a network were to charge Acme Video, a hypothetical video-on-demand provider, a higher price for network access, it might do so not because Acme is Acme or because Acme provides video-on-demand, but because video-on-demand uses a lot of bandwidth and Acme happens to be a well-known provider that is easy to identify. On the other hand, the network operator might charge Acme the higher price either because

---

"free riding" on network infrastructure because they make big profits in which network owners do not share. The argument is weak. Applications providers are no more free riding on network platforms than vice versa. Consumers do not purchase Internet access from network operators just to cruise the network; they subscribe to reach on-line content and services. Just as network operators do not share in the profits of such applications providers, nor do they share their subscription revenues with the applications providers that consumers pay to reach. Moreover, it should be noted that many applications providers fail, and while network operators may not share in the profits of the successful ones, nor do they share the investment risk and losses from applications ventures that fail. What may look like free riding to the platforms may look like portfolio skimming from the other side.

Acme happens to be a rival in a particular complementary market or because the network operator knows Acme has deep pockets and will pay a lot not to have its traffic consigned to a slow lane. As discussed above, these latter two motivations may have little to do with cost recovery and carry some risk of anticompetitive harm or other allocative inefficiency.

Discrimination targeted at particular uses is potentially more neutral, although it is not necessarily better than discrimination by user. If higher prices are charged only based on whether a particular use is one that competes with a business of the network, then it may be anticompetitive. For discrimination by use to be better than discrimination by user, the categories must be chosen because they are reasonable proxies for costs imposed on the network rather than proxies for competition.

The most neutral of the three options for price discrimination is usage-based pricing, *i.e.*, charging for the amount of traffic an applications provider does or expects to put on the network. Some forms of usage-based pricing blur the line between targeted and non-targeted price discrimination. For example, if a network operator were to meter traffic and, as congestion developed, turn some capacity into a priority "lane" that any user could select for a fee, then the pricing would be non-targeted. If, however, the network operator mandated increasing fees as an applications provider crossed progressively higher thresholds of traffic volume, then the price discrimination would be targeting such high-volume users for higher access prices.

The most risky forms of price discrimination for competition and innovation, therefore, appear to be those where the network operator can target particular uses or users for higher prices. The least risky forms of price discrimination are those that charge for priority on a usage basis, where each applications provider can decide whether to purchase priority. While it may still be possible for pricing mechanisms to be designed to coerce particular applications providers to pay more, a posted menu of prices for priority based on usage raises many fewer concerns than targeted pricing based on use or user.

The costs and benefits of price discrimination by networks to applications providers thus vary with two sets of factors: the motivation for price discrimination and the method by which it is accomplished. Charging for priority in the presence of capacity constraints and congestion is more likely to yield benefits than is selling priority in the absence of capacity constraints. The first can represent an efficient response to scarcity; the second runs the greater risk of being an inefficient exercise of market power. Next, charging for priority based solely on usage rather than setting terms that target particular uses or users is more likely to avoid anticompetitive uses of price discrimination.

A basic taxonomy of network price discrimination, compared by level of anticompetitive risk, is summarized in the table below.

A Simple Taxonomy of Price Discrimination by Networks

	Targeted Pricing	**Non-Targeted Pricing**
Priority with Capacity Constraint	Moderate anticompetitive risk (??)	Lowest anticompetitive risk (best option)
Priority without Capacity Constraint	Highest anticompetitive risk (worst option)	Moderate anticompetitive risk (??)

The schema presented above suggests that not all discrimination need be equally harmful, in turn implying that the costs and benefits of network neutrality regulation will differ depending upon which kind of conduct it prohibits. To the extent there can be benefits to price discrimination itself, prohibiting even the comparably benign forms of discrimination might forego benefits in return for the prevention of less substantial harms. The next section addresses the implications of this possibility for regulatory policy.

C. Conclusion: Policy Alternatives Going Forward

The different motivations and methods of price discrimination raise the possibility of policy solutions that focus selectively on the most harmful kinds of discrimination without prohibiting other non-neutral conduct that could yield net benefits. Policy could regulate actions most likely to foreclose competition either among applications providers or between applications providers and the underlying network. Such regulation would not need to preemptively prohibit networks from offering a non-targeted menu of access tiers available to all applications providers regardless of their identity or type of service. This more selective focus is consistent with two commentators' recommendation for regulation that precludes network owners from discriminating among data packets routed on their networks based on the identity of users or uses.[30] It reduces the risks of targeted discrimination without banning discrimination altogether, thereby preserving some of the potential

30. *See* Brett Frischmann & Barbara van Schewick, *Yoo's Frame and What It Ignores: Network Neutrality and the Economics of an Information Superhighway*, 47 JURIMETRICS J. (forthcoming Summer 2007), *available at* http://ssrn.com/abstract=1014691.

benefits of upstream price discrimination by network operators. In terms of the chart displayed above, regulation would rule out the two left-hand quadrants. One might also try to rule out the lower right hand quadrant because the priority there is discretionary rather than driven by physical capacity constraints. Capacity constraints may be hard to observe and monitor, however, so regulation might as a practical matter do better to focus more on the method (*i.e.*, pricing structure) than on the motivation (*i.e.*, existence or not of real capacity constraint) for price discrimination.

There are different ways in which departures from non-targeted pricing, and the associated hazards for competition, could be regulated. One alternative is to have a basic rule that prohibits outright blocking of any (legal) applications provider, coupled with a regime of *ex post* enforcement against price discrimination that can be demonstrated to be anticompetitive. The approach here is primarily an antitrust-style approach. It has the virtue of not prohibiting much conduct in advance of proven anticompetitive effects, but would involve the courts and enforcement agency in assessing the detailed terms of each individual deal that came before them. The no-blocking rule would mean that such an *ex post* regime would differ from general U.S. antitrust law, which generally does not prohibit outright refusals to deal.[31] The focus of *ex post* enforcement would more likely be on whether the terms of trade were anticompetitive or not.

An alternative solution would be to impose some *ex ante* restraints on those terms of trade through a network-neutrality rule that imposes a light form of common carriage on the network operator. A modest rule might still allow networks to offer different access terms to applications providers but would require that those terms be transparent and available to all such providers. One promising proposal combines such an approach with *ex post* enforcement against any anticompetitive uses of price discrimination by a network.[32] The devil is likely to be in the details for either of these approaches, and detailed exploration is beyond the scope of this brief essay. The important point is that intermediate solutions exist that can dampen the worst potential harms of network access discrimination, without altogether banning all price-mediated prioritization of network traffic. In light of the open questions that each side of the debate raises, such intermediate solutions warrant further development.

Finally, the most essential long-run strategy to reduce the risks of anticompetitive discrimination raised by the advocates of network neutrality is to focus on horizontal competition rather than vertical

31. *See* Verizon Commc'ns Inc. v. Law Offices of Curtis V. Trinko, LLP, 540 U.S. 398, 408-09 (2004).

32. ATKINSON & WEISER, *supra* note 4, at 12.

regulation. If the competitive progress of the U.S. telecommunications market can be maintained through effective radio-spectrum policy, network interconnection rules, and vigilant antitrust (particularly merger) enforcement, then network neutrality concerns will diminish. Congress and the FCC should, therefore, not lose track of longer-term structural solutions for improving competition and innovation in the broadband market, and should ensure that any interim regulation they impose will not remain in force as market conditions no longer justify them.