

Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets

Andrew Odlyzko

School of Mathematics and Digital Technology Center
University of Minnesota
127 Vincent Hall, 206 Church St. SE
Minneapolis, MN 55455, USA
odlyzko@umn.edu
<http://www.dtc.umn.edu/~odlyzko>
Revised version, January 17, 2009

Abstract. Network neutrality as such may fade from public interest discussions, but historical precedents going back for centuries argue that the underlying issues will continue to be debated. Those issues revolve around the basic tension between efficiency and fairness in markets, a tension that has never been completely resolved. Further, this tension is unlikely to ever be resolved, since our well-documented inability to predict the development of technology and its impact on society mean that no fixed set of rules can work indefinitely.

Should net neutrality or some similar set of rules come to dominate (either because of market forces, or through regulation), attention would likely turn to other parts of the economy that might be perceived as choke points for social, economic, and political activities. If Net search becomes as important as the Google stock price seems to imply, for example, it might be the focal point for such concerns.

Future controversies are of course matters of speculation. On the other hand, net neutrality and its close relatives, such as common carriage for the Internet, are current issues that have to be decided soon. What appears to be missing in the current debate is a discussion supported by reliable data of the basic fundamental economic question, namely whether a net neutral communication infrastructure can be viable. And if it is not, just how far from net neutrality is it necessary to move? Should pricing for Internet access be dependent on the income of the user, for example? There are certainly precedents for such extreme policies. However, they usually occurred in cases where there was a need for very high revenues for service providers.

There are arguments that a net neutral communication infrastructure should be viable. But to get there would require a major restructuring of the industry. The prospects of that are tied up not just with politics, but also with some of the great paradoxes of the current financial markets.

It is possible to argue that the best outcome might be to have Google defeat AT&T in the battle over net neutrality, but then (and likely in any case) society might have to get ready to regulate Google!

1 Introduction

As Yogi Berra and others are supposed to have said, “It’s tough to make predictions, especially about the future.” This definitely applies to society’s reaction to technology. History is full of mistaken predictions. As just one example, in 1998 Sergey Brin and Larry Page [5] claimed that

The goals of the advertising business model do not always correspond to providing quality search to users. ... we expect that advertising funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers. ... we believe the issue of advertising causes enough mixed incentives that it is crucial to have a competitive search engine that is transparent and in the academic realm.

However, Messrs. Brin and Page realized their mistake, and went on to achieve fame and wealth building Google, which (at least so far) has managed to be just what they had predicted was impossible, an extraordinarily successful search engine that is supported by ads, and is both proprietary and extraordinarily opaque. This is a common pattern in the history of technology, that even the inventors and promoters of an innovation have mistaken notions of the impact it will have on society, yet often they do succeed in spite of this. As another example of this phenomenon, in the early years of railroads, it was widely expected that their spread would severely cut back on horse populations. Instead, railroads stimulated horse breeding, since rails were great for transport once one got to them, but the “first mile” problem of reaching those rails required horses!

While detailed predictions appear impossible most of the time, there are certain persistent trends in history. One of these trends is that some mistakes keep getting made again and again, in spite of all the experiences of previous centuries. Telecom seems especially prone to this tendency as is discussed below, which makes it possible to predict that many ventures will fail (an exception to the general rule about difficulty of prediction), but complicates planning in other ways. A related persistent trend is that controversies can frequently be classified into certain patterns that recur over and over again. This trend, which applies to the net neutrality debate, strongly suggests that the basic issue involved in it is not going to disappear. The issue is how much control service providers should have over the bits that society relies on so much.

The net neutrality issue has been pronounced dead and buried several times, but it has bounced back, typically when a service provider said or did something that struck some people as outrageous. Thus at the end of 2005, Ed Whitacre of SBC (now renamed AT&T), when asked about business prospects of companies like Google, responded [42]:

How do you think they’re going to get to customers? Through a broadband pipe. Cable companies have them. We have them. Now what they 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.

This led to extensive protests and discussions that appeared to peter out after a while. However, at the end of 2007, Comcast stirred up a new round of controversy by spoofing packets to disrupt peer-to-peer (P2P) transmissions of its customers.

The net neutrality debate is often pictured as a contest between the two most prominent corporate champions of the opposing sides, AT&T and Google. But the underlying issue predates both companies by centuries. It was never resolved completely, since it arises from a conflict between society's drives for economic efficiency and for fairness. There is no reason to expect that this conflict will lessen, and instead there are arguments that suggest it will intensify. Should something like net neutrality prevail, the conflict would likely move to a different level. That level might become search neutrality. (And allegations about discriminatory behavior of a web search provider have surfaced recently in China [39].) Or, to take another currently popular concept, if "cloud computing" does become as significant as its enthusiasts claims, it could lead to dominance of a single service provider. The effective monopoly of that dominant player could then become perceived as far more insidious than any of the "walled gardens" or "intelligent networks" that telcos would like to build.

This brief note outlines some of the issues related to net neutrality, and the controversies that might replace it. It concludes by raising a key question that appears missing from the current net neutrality debates. It is that of a quantitative discussion of just how much it costs to build out and run a broadband network, whether such a network can be viable under net neutrality, and if not, how far away from net neutrality should one expect (or allow) service providers to depart.

What is acknowledged, whether explicitly or implicitly, in most discussions of net neutrality is that the basic issue is of price discrimination. There are frequent claims about the need to manage network traffic, but when one gets deeper into them, one typically finds complaints about "5% of the users generating 50% of the traffic" and the like, which have less to do with providing adequate service, and more with a way to apportion the cost. And the usual practices that have been observed have tended to be more about restricting some applications than about managing congestion, say. For example, consider the Comcast vs. BitTorrent affair that surfaced in 2007. The FCC in its order to Comcast to end its practices [12] noted that Comcast was throttling BitTorrent traffic even when there was no congestion, while allowing heavy users to carry out unimpeded even during peak congestion periods, as long as they were not using BitTorrent. And Comcast changed its story several times, when its previous explanations were found to be at variance with facts. Now it is true that this FCC decision was by a 3 to 2 vote, and one of the opponents, Commissioner McDowell, noted in his dissent that the evidence was "thin and conflicting," with further complaints about lack of signed statements and the like. So perhaps the FCC decision was wrong in its statement of facts. But the description in [12] is typical of what one finds in history. Companies engaging in discriminatory practices try to conceal it, dissemble when caught, and cite various seemingly less objectionable objectives (safety, performance, ...) as the reasons for their policies. Very seldom do they come out and admit what they are doing.

It should be emphasized that from the standard economic point of view, there is nothing nefarious about attempting to introduce differential charging. It's just that social acceptability of such practices is at wide variance with the recommendations of the conventional economic models. There are many precedents for price discrimination in all commerce, and in transportation and communication in particular. With more control, service providers can obtain more funding, which may be used to build out networks. But in the context of net neutrality, this issue is not discussed very extensively. And when it is, it is treated in a curiously one-side manner, just about discriminating among content providers, and without specifying just how much discrimination is needed. Service providers argue that they can finance the buildout of broadband networks only by charging extra fees to parties such as Google. But so far the main change that has taken place in the U.S. (aside from a few providers interfering with P2P transmissions) is that in some states telephone companies have won the right to deploy broadband only in those neighborhoods they choose. (There are some provisions about not discriminating against the poor, but those are vague and almost surely unenforceable. Traditionally, both cable and telephone companies had obtained the right to string their wires in municipalities by promising to serve every one in each area equally. But that was not always so in the early days of the telephone industry, see [14].) If broadband networks are horrendously expensive to build, as is often implied, should one perhaps go further, and make the price depend on the income of the user? That's how a Harvard education is priced, after all. There is a "list price" of about \$45,000 per year, and families with incomes up to \$180,000 per year (as announced at year-end 2007) can obtain discounts by giving up their financial privacy and submitting detailed accounting of their income and assets. So we might price basic Internet connectivity at \$1,000 per month, with discounts for those who can't afford that much and are willing to reveal detailed financial information. Such a policy could also be accompanied by additional charges for every bit that is actually transmitted, with the charge determined by the value of that particular transaction.

Comparing Internet access to college tuition may be extreme, but it helps frame the discussion by exhibiting the whole range of network access and pricing policies. There is a range of net neutrality choices at one end (depending on the definition of net neutrality one chooses) and a Harvard-like policy at the other. Which policy will be sufficient to produce the networks that will be needed?

Contrary to many claims of opponents of net neutrality, networks are not very expensive to build (as will be discussed later), and some simple calculations suggest that a net neutral communications infrastructure could be viable economically. But such an infrastructure might enable even more extreme forms of price discrimination by players such as Google, and might then lead to new controversies and demands for new forms of regulation.

2 The value of bits and content versus connectivity

It helps to consider how much value is attached to different types of bits. Now value is hard to define (there is market value, different in different markets, and use value), so as a proxy, let us consider how much users pay for various telecommunication services. We

Table 1. Value of bits: Price per megabyte of various services.

| service | revenue per MB |
|----------------------|-------------------|
| wireless texting | \$1000.00 |
| wireless voice | 1.00 |
| wireline voice | 0.10 |
| residential Internet | 0.01 |
| backbone Internet | 0.0001 |

then come up with a listing that looks something like Table 1. These are just prices paid, and so do reflect varying costs, competition, and pricing policies in various degrees.

The numbers in this table should not be taken as too precise. They were chosen to be nice round numbers that are in approximately the right ranges. For example, it appears that SMS messages might be something like 30 bytes on average, and in much of the world they cost about \$0.10 each, which results in a price of \$3,000 per MB, but that is close to the \$1,000 listed. Many of these services, especially in the U.S., are sold in bundles (a topic to be covered in much more detail below), so per-unit or per-MB pricing is not visible to the users. But generally speaking a residential user who transmits 3 GB/month and pays \$30/month for a broadband connection (figures that seem typical for some countries) is effectively paying about one cent per MB. Similarly, a small ISP that has to purchase access to the backbones from larger service providers and pays \$20 per Mbps per month (which would be a high price in the U.S. as this paper is being written, but applies in some parts of the world) on the 95-percentile basis, with a 1.5 peak-to-average traffic ratio, is paying about one hundredths of a penny per MB. Thus all the numbers in the table can be defended with data as approximately correct, usually within factors of two or three. It is the great, orders of magnitude, disparities between them that are of interest.

Price per bit is just one measure of a communication service. One should also keep in mind profits, since that is what drives investments, and total revenues. It appears that nothing in telecom is at the profit levels of wireless texting (aside possibly from ring tone downloads, or services such as Caller ID). This helps explain why the telcos are attracted to “walled garden” approaches, where they hope to obtain similarly high profits from narrowly defined services. Note also that the basic search/ad services of Google generate very little traffic, but are extremely profitable, leading to the “Google envy” that appears to be driving so much of current telecom planning. But in terms of revenues, an estimated 70-80% of worldwide telecom revenues come from voice [19] (with wireless increasingly dominating). And the bulk of the remainder is from other connectivity services (such as texting, or, to the extent one can allocate revenues from Internet access, from email). In particular, this all by itself suggests that the claims of a high level AT&T manager [44] that net neutrality “is about streaming movies” is mistaken. If anything, net neutrality and the future of telecom in general, is far more about voice and the services that will succeed voice in importance. In particular, the key issue is how to extract payment from customers when voice, the most valuable service they receive, is a free feature of the basic broadband connectivity they buy.

The comment that net neutrality “is about streaming movies” is a reflection of two fundamental and very deeply ingrained false myths. One of them is that “content is king.” This was the key motivator behind European service providers throwing away about 100 billion Euros (almost \$150 billion at today’s exchange rates) on 3G spectrum auctions, and then more on top of that for misdirected technology development and deployment. The promise was that wireless “content” services (movies, music, location based services, and the like) would double average monthly revenues. Yet even a superficial look at history would have shown the “content is king” myth to be false [29, 30]. Connectivity has always been valued far more than content. This is slowly and painfully being rediscovered here and there. For example, in the briefing by Takeshi Natsuno, “one of the principal architects behind DoCoMos wildly successful 1999 launch of i-mode,” where “one message became abundantly clear: content is not king” [40]. Although wireless data revenues are growing, they come overwhelmingly from texting, a connectivity service, and that is what users care about the most [2]. For another, more recent, scholarly paper that provides broad historical and statistical support for the observation about the limited role of “content,” see [1]. Still, the “content is king” myth has reigned for over two centuries [29, 30] in spite of extensive evidence against it. And it reigns not just among business executives and policy makers, but among scholars. The recent literature on net neutrality, for example, relies on models of transportation networks as content delivery mechanisms (see, for example, [10, 36, 38] and many others), and it is only rarely that one hears even a brief parenthetical mention of a dissenting voice (as in [18]). So it seems a safe bet that the “content is king” myth will guide decision making in the future. Hence many decisions will be made in predictably incorrect ways, in particular about regulating net neutrality or about trying to exploit its lack.

The other false myth embedded in the “streaming movies” quote is that movies should be delivered in real-time streaming mode. Video is already a very large fraction of Internet traffic. And it appears that an increasing fraction is being streamed. But there is little reason, except for ingrained prejudices, why that should be so. With very few exceptions (voice conversations and videoconferencing the primary ones), it is far more effective (in terms of providing quality service, and in keeping the network manageable and inexpensive) to transfer segments of movies or music faster than real-time for replay from a local buffer [28, 35]. But this myth is also deeply ingrained, and seems impossible to dislodge. So it is likely to continue misleading both service providers and policy makers. Service providers will build out expensive networks designed for streaming video that will waste their shareholders’ and their customers’ money. And policy makers will worry about how to regulate such features.

The combination of deeply embedded but false myths about communications, combined with general inability to predict both developments in technology and society’s reaction to technology, mean that no fixed set of rules is likely to be satisfactory, as reality will be intruding and making the best laid plans irrelevant. But some of the issues that will be controversial seem very easy to predict, in particular the degree of price discrimination that service providers are allowed to practice.

3 Price discrimination precedents

The wide disparities in prices seen in Table 1 help explain the strong incentives towards price discrimination that exist. Similar incentives have existed for ages, and they frequently led to differential pricing, both in communications (see, for example, [31] or the more detailed discussion in [29]) and in transportation [34]. As just one example, let us note that in 1829 the toll schedule on the Stockton and Darlington Railway, the pioneering locomotive-powered public railway, had charges of 0.5 (old English) pence per ton per mile for coal destined for shipment out of the region, 4 pence per ton per mile for coal for use in the region, and 6 pence per ton per mile for general merchandise (Table 5 on p. 83 of [20]). This is just one of myriad examples one can cite (for others, see [21, 34], for example). But it does illustrate several relevant points. One is that charges were dependent not just on the nature of the cargo, but also on its destination, so we had something similar to today’s much disliked and feared “deep packet inspection.” (And there were attempts to conceal nature of some shipments, with penalties imposed when detected.) Another is that such charges were not the result of simple economic optimization, but often involved heavy politics. In particular, the low charges for coal destined for shipment were the result of coal mine owners in a different region trying to avoid competition from the coal mines that the Stockton and Darlington Railway was to serve (see [20] for details). They persuaded Parliament to impose what they thought was an unprofitable rate that would ruin this line. And yet another lesson (about difficulty in forecasting technology) is that this attempt failed, since those coal mine owners did not anticipate progress in railway development correctly, and the supposedly ruinous rate turned out to work quite well. Finally, yet another lesson is that although the Stockton and Darlington tolls differed substantially according to the nature of freight and whether it was destined for shipment or for in-region use, they were otherwise uniform, not distinguishing between shippers, or locations, being based strictly on mileage. This represented the other major factor to be discussed later, namely that price discrimination has traditionally been disliked and feared, and has usually been strictly limited.

In view of the historical precedents for price discrimination in communication and transportation, some of the arguments of net neutrality proponents are questionable. For example, Vint Cerf of Google [25] has claimed that since both residential users and companies like Google pay for access to the Internet, it would be unfair for “access providers ... to step in the middle and create a toll road to limit customers’ ability to get access to services of their choice.” Yet such tolls have been frequent, and have often been used to shift costs of the entire communication system around. Note that Table 1 shows that a communication between a residential user and Google costs about 100 times as much to the user as to Google. (The differential is likely much more than 100, since Google may very well be getting free connections to some access providers using its own fiber or wavelengths. Such arrangements probably cost much less than the transit fees that were used in computing the numbers in that table.) But in the past, in the traditional voice telephone network, business customers frequently paid far more than residential ones. Sometimes this was voluntary (as in the provision of toll-free numbers), but most of the time especially high charges for businesses were imposed on them. This started out with the very first

telephone service. The private company that controlled the Bell invention used its patent-protected monopoly to simply demand twice as much for business uses of the telephone as for personal use [29]. This policy of charging business customers extra continued under regulation. As a rough approximation, it appears that back about half a century ago, business customers in the U.S. paid about two-thirds of the cost of the telephone network, and residential users for one-third. Over the last decade, the balance appears to have shifted, and so corporations are escaping the burden they used to bear. (Whether that is good or not is another issue.)

Thus from a historical perspective, there is nothing novel about service providers attempting to erect toll booths and charge extra for traffic to and from Google. The marketplace is full of various indirect payment arrangements. In principle it would not be any more unusual for Google to pay access providers to allow residential customers to connect to the whole Internet (and not just to Google) than it is for advertisers to pay Google to provide search results to end users who in most cases do not pay anything to those advertisers.

4 Privacy and price discrimination

The incentives for differential charging are well understood, but seldom discussed in public. But they are discussed in private. As one example, an antitrust lawsuit made public an email [43] of Aug. 21, 1997 from Warren Buffett, "the sage of Omaha," to Jeff Raikes of Microsoft, in which Buffett wrote that

[Alexander Graham] Bell should have anticipated Bill [Gates] and let someone else put in the phone infrastructure while he collected by the minute and distance (and even importance of the call if he could have figured a wait [sic] to monitor it) in perpetuity.

Alexander Graham Bell did not collect by the importance of a phone call. However, he, or rather the company that he helped start with his patents, certainly did charge by the minute and distance (which has traditionally correlated with value and willingness to pay, the rationale for distance-sensitive charges in postal systems before the Penny Post Reform). And history does record instances of telephone companies attempting to restrict usage of hotel phones to hotel employees. There was even a telco protest against the use of a telephone to alert the fire department to a fire on a neighbor's property on the grounds that the phone was for personal use only [29]. So the incentives to price discriminate have been understood for a long time, and were frequently pursued to the extent (and often beyond that) allowed by technology and law. Technology was often a limiting factor, as a rich man could dress in rags when showing up for a train or coach ride, and could send a servant to buy goods at low prices. Further, rules and regulations also often interfered with what sellers might have wanted to do.

The information and communication technologies revolution, and the Internet in particular, are leading to unprecedented opportunities for differential charging. Buyers' willingness to pay can be estimated from detailed knowledge of their behavior, and their usage

can be monitored and restricted, thus inhibiting resale, which has been a traditional escape route from extreme forms of differential pricing. The thesis proposed in [26, 32] is that the growing erosion of privacy can be explained by companies preparing to exploit the growing opportunities for price discrimination. But employing price discrimination is like playing with fire, as is shown by past experience.

5 Fair and efficient markets

There are strong incentives for fine-scale differential pricing, with the Holy Grail of commerce a system in which every item or service component is priced at the buyer's maximal willingness to pay. But there are also several countervailing tendencies. They include the generic benefits of bundling, which is how flat rate services can be viewed. (One of the great ironies of the communications industry is that one segment fights against attempts to make it unbundle cable TV channel. At the same time, most of the industry keeps complaining that flat rates are unsustainable on the Internet and elsewhere, and supports development and deployment of technologies such as IMS that could make fine-grained charging feasible. And yet all this time flat rates are spreading!) But in addition there is the willingness of users to pay more for flat rates, as well as the substantial effect that flat rates have in stimulating usage. These factors are described in detail in [13, 23, 29, 31]. They frequently swamp the incentives towards fine-scale discriminatory pricing, in what may be perceived as another triumph of behavioral economics over the standard economic models. It appears likely that for relatively inexpensive goods and services that are used frequently, flat rates or other simple pricing plans are often of benefit to sellers. This issue won't be considered further here, except to note that these effects have been visible in communications at least since the Penny Post Reform in Britain in 1840, and that the telecom industry has refused to learn from them. Hence it is safe to predict that service provider policies as well as regulation will often be misdirected through ignorance of such effects.

The main constraint on price discrimination comes from society's dislike of the practice. This dislike is reflected in customs, laws, and regulations. Society has struggled for centuries with the conflicting incentives, one to allow differential pricing in order to stimulate greater production, and the other the desire to limit it. For many segments of the economy, essentially untrammelled markets were allowed to operate, subject to minor antitrust and related constraints, and the discrimination that arose there was tolerated. But for key communication and transportation industries that was not regarded as adequate, and common carriage was the doctrine that was applied. Common carriage served to limit discrimination, but, in a nod to economic incentives, did not prohibit discrimination entirely. For example, senior citizen or student discounts were and continue to be accepted. What was illegal were "undue or unreasonable" preferences, vague terms that provided lifetime employment to legions of lawyers (and, in more recent times, economists). At this moment, in the U.S., common carriage has been waived for Internet services. But it could easily come back. And even more stringent constraints could come. The interesting comparison is with that great revolutionary technology of the 19th century, the railroads. They were subject to common carriage rules from the beginning. Not only that, but their early charters had embedded in them structural separation (similar to the one that prevailed

on canals). The railroad company was expected to provide a rail road, and have carriers use their own wagons and motive power. But this quickly broke down. In the words of a British observer writing soon after the change took place (p. 525 of [22]):

... the privileges and rights contemplated, as well by the companies as by Parliament, were merely those necessary to enable them to construct and maintain a road, which was to be open to all who might desire to use it, on the payment of a certain toll to the company. ... A colossal monopoly, never contemplated by Parliament, nor foreseen by the companies themselves, had come into being.

But even when vertical integration was conceded to railways, they were constrained to charges that depended only on mileage travelled, class of service, and (for freight) nature of good carried. Thus an enthusiastic supporter of railways, and an advocate for them in the small package wars (described briefly in [34]) claimed in 1847 (p. 12 of [3]), after a discussion of how valuable differential pricing would be for a particular railways, that “The people of Kent will never rest satisfied with, the public will not tolerate, the legislature would grant a remedy for the correction of a [scheme with high per-mile fares on some routes and low fares on others].” So even though the economic incentives pointed towards differential pricing according to route, and those would have helped to spread the network more widely, this observer was not willing to go that far, although he was supportive of differing charges for different classes of passenger service and for different types of freight.

The vertical integration followed by railroads in getting into carriage of goods and passengers was usually justified on the grounds of safety. Railroad managers complained they could not protect life and limb unless they controlled the entire operation on their line. And they did have evidence to support their contention, just as today’s telecom service providers argue they need fine scale control over the packet traffic in order to prevent malicious attacks, assure customers of decent connections, etc. But it appears that much of the motivation for the vertical integration of railroads came from the the greater profit opportunities that railroads saw in operating that way. This is shown both by the occasional quotes from managers, and the fact that railroads did manage to operate their trains over the lines of other railroads with which they had friendly business relationships. (This was an early instance of the phenomenon demonstrated to hold in telecom a century and a half later, that in the absence of clear boundary, it is hard to regulate a complicated system if the main operators do not wish to cooperate [11].) What ensued was a period of intense development of railroads, and also of extensive price discrimination that aroused strong protests. In the words of a report from a committee of the U.S. Senate in 1874 (cited on pp. 382–383 of [17]), what was desperately needed was

a remedy for the evils of unjust discrimination against one locality in favor of another, or in favor of one description of trade at the expense of another ... and of uncertainty and favoritism by means of special contracts, rebates, drawbacks, and the thousand and one other means by which a rich and powerful company may, by the secret adjustment of rates, impose upon the public.

The complaints were not just about level of railroad charges (which were declining quite rapidly during the 1870s). The concern was more about the fairness of the system. A contemporary comment on the situation was [16]:

But the fact that the charges are so low does not make *differences* in charge bear any less severely upon business. A difference of five cents per bushel in the charge for transporting wheat a thousand miles is a small matter, taken by itself. It would be weeks before it would make a difference of one cent to the individual consumer of bread. But if a railroad makes this reduction for one miller, and not another, it will be enough to drive the latter out of business.

The dislike of price discrimination has deep and still poorly understood roots (see [32] for some references) at the level of individuals. Even people who receive discounts are often suspicious as to whether they are being treated fairly, for example. But there is also a deeper concern that arises when differential pricing becomes widespread and remains opaque, a concern that affects the moral foundations of a free market. In the case of 19th century railroads, society, led by business groups, demanded a more transparent and level playing field. The setting where a monopoly infrastructure business, in pursuit of its own ends, could take arbitrary steps that would ruin one business and make another succeed, were regarded as inimical to a really free market. It resembled far too much the widely disliked markets without property rights, dominated by a capricious political power. So what followed was a long period of increasingly stringent regulation. Critics later complained about “regulatory capture,” in which the regulated industries controlled the regulated. And that was certainly true. Those industries did gain stability, while society gained a precarious sense of fairness. (It was precarious since the differential pricing incentives could not be ignored completely, and all those legions of lawyers and economists argued about a myriad of cases, as various players tried to push the boundaries of acceptable behavior.) Some economic efficiency was lost, at least in the short term, but that seemed to be regarded as acceptable.

6 Future of pricing

The history of railroads suggests some general trends in how pricing will develop [32]. One is that the differential pricing incentive will continue to operate, and that sellers will continue pressing against, and often transgressing, whatever rules are set up. (Railroad history is full of examples of companies pledging to refrain from price discrimination, either because of new laws, or because of voluntary cartel arrangements with their competitors, and then immediately breaking such pledges.) The eager reception that technologies such as Deep Packet Inspection, IMS, and others receive from service operators surely reflects this tendency.

To what extent such technologies are actually going to be used is another question. There will be much experimentation to see what society will accept. The Harvard-like policy might be the dream for all sellers, but will seldom be achieved, and certainly not in the near future. Starbucks will not suddenly declare that the price of a cup of coffee will be \$2, \$5, or \$7, depending on how much Starbucks thinks a particular individual might be able and willing to pay, based on that person’s history of Starbucks purchases as well

as income and assets, with a “standard list price” of \$10 per cup for those who insist on remaining anonymous. (But note that prescription drugs in the U.S. are approaching that type of pricing.) On the other hand, Starbucks (and other coffee shops, there is nothing special about Starbucks) will likely experiment with what is nowadays called “behavioral targeting,” in which individuals get customized ads depending on a variety of factors. The ones that are usually cited are time of day and proximity to the nearest store. But the incentive will clearly be to incorporate into the creation of these ads particular individuals’ willingness and ability to pay. So, for example, a lover of Starbucks coffee who is not price sensitive might receive on his cell phone discount coupons that average \$0.50 per cup, while somebody who does not care whether he drinks coffee from Starbucks or McDonald’s might get discount offers of \$1.50. The offers will vary from instance to instance, so they will not always be \$0.50 and \$1.50 for those individuals, making it harder to detect what is happening. But the average discounts will be different, and will enable sellers to realize some of the benefits of the ideal form of differential pricing.

This may sound like a “brave new world,” but glimpses of it do occasionally appear. For example, consider the descriptions of a car dealer network in [15] that does not post prices, and determines what to offer to its (poor and financially unsophisticated) customers based on detailed information about their position. And, of course, we do have Harvard and other elite private schools that practice extreme forms of differential pricing, and medical care appears to be moving in that direction [32]. We do not know exactly what forms of price discrimination society will accept. So we should expect experimentation, hidden as much as sellers can manage, but occasionally erupting in protests, and those protests leading to sellers pulling back, at least partially. And occasionally we should expect government action, when the protests grow severe. (However, based on historical precedents we should expect governments to often pay lip service to public demands, while in practice tolerating and even promoting differential pricing, because of its positive economic effects, see [32].)

7 How to justify Google’s stock price

The precise form that price discrimination will take is impossible to predict, and the Starbucks scenario above is just one way the future can be approached. But there is huge potential in it, and it may be what has led to the huge stock market valuation of Google. Serving up the standard online ads has turned out to be extremely profitable. (In another instance of the difficulty of technological predictions, this is contrary to general expectations as recently as the peak of the Internet bubble in 2000, when the consensus was that such ads were not a promising revenue source.) And there are plenty more ads that can migrate online. But there are limits, since the advertising field is not all that large. (The total amount spent on advertising is lower than spending on voice telephony.) But if “behavioral targeting” can be exploited, extensive price discrimination would open up entirely new, and much larger, revenue sources. (Enterprises already spend on direct marketing and other promotions amounts comparable to what they spend on ads, and effective differential pricing could replace and augment that.) And the company in control of the process, the one with the information about customers, and in control of the delivery of the offers, could

potentially end up in control of product pricing, essentially relegating the good and service providers to a commodity role. The profit potential of such a role is gigantic.

Yet another way that Google could justify its high valuation would be if the current hype about “cloud computing” becomes reality. If most of our data and computing migrate into a distributed network of processors and storage, entirely new opportunities would arise. Given the economics and the technology, with the “first mover advantage,” “economies of scale,” and similar factors, cloud computing could easily become dominated by one or a handful of players, who might then have huge power, and a corresponding profit potential. And Google is ideally positioned to grab such a role, since its core expertise is in deploying a distributed information technology infrastructure. In either scenario, though, Google (or Yahoo! or Microsoft, or Facebook, or some other entity that might reach a position to exploit its position in the ways outlined here) would have to tread very carefully, since the public could easily decide its practices amounted to “being evil.”

For Google (or some other entity) to fully exploit the potential of differential pricing or of cloud computing (and the latter would surely also involve extensive price discrimination), it is necessary to have a communications infrastructure that either that entity controls, or that is approximately net neutral. Otherwise, the communications network could grab the profits by charging differential fees to Google that would absorb most of the benefits, or could deploy its own competing “cloud” which might be less efficient, but could exclude competitors.

8 Fairness and rules

The basic conclusion is that whether AT&T or Google wins the net neutrality battle, the outcome at a high level may be similar, namely society exposed to the prospect of an unprecedented degree of discrimination. It is doubtful that competition could mitigate the risks. On one hand, there is doubt whether one could have a viable system with enough separate physical networks for effective competition. We already have, in most of the United States, two separate networks. The case is made below that there are already sufficient revenues to allow both networks to deliver high speed connectivity. But a duopoly can hardly be expected to provide a competitive market. On the other hand, it does not seem feasible to pay for five separate physical networks, say. (And even if one could get enough funding, one would then have to face the issue that surfaced prominently with railroads about 170 years ago, namely whether such wasteful duplication would be preferable to the inefficiency of regulation or government ownership.)

It is not even clear whether a duopoly would be stable. (And, it must be admitted, that it is not even certain that the wireline networks will continue to maintain their revenues, since it is voice connectivity that is the most valuable service, and voice is rapidly migrating to wireless.) In cyberspace, the “winner take all” phenomenon is very pronounced. And in addition, in the online world, strategic gaming and tacit collusion are easy (as the electric power “crisis” in California demonstrated recently). Furthermore, we have extensive historical evidence that competition often leads to increases in price discrimination.

To what extent the winner of the net neutrality debate might actually be able to exploit the opportunities that victory would offer will depend on how people react. It is not

just the existing laws and regulations that matter, but public reaction. Because one of the key lessons from the past is that governments do respond to popular concerns about the functioning of markets. People have often called on governments to intervene when they got upset with what they perceived as unfair rules. (Whether that intervention takes place through laws, regulations, or courts intervening to invoke common carriage rules is a different question, and will surely differ from country to country, as it has in the past.) We see this happening today in the European mandate to lower international roaming fees in wireless voice [4]. On purely economic grounds those high fees can easily be justified since they are charged on the calls that are likely to be most valuable. (And that those high discriminatory fees existed even in a competitive market shows something that has been visible for a long time, namely that competition is not necessarily a barrier to price discrimination.) But they offended the public sense of fairness, and so they are being reduced through explicit government mandates and the threat of such mandates. People did not resort to the technical means of circumventing those fees, but instead insisted on changing the rules of the marketplace. Hence we should not take the general lack of concern about their privacy by people as indicating that any uses of private information will be acceptable. If people perceive such information is used against them (as in differential pricing, which typically elicits strongly negative reactions), they are likely to call on government to fix the problem.

The general conclusion is that some form of government intervention, to set the rules, is inevitable. (And at some point it may be welcomed by the players, just as government intervention was welcomed in the end by the railroads.) Society needs basic rules to operate by, and modern technology creates potential scenarios that old rules did not cover. But we need to remember also that it is not easy to regulate markets, especially ones in cyberspace, and especially when policy makers labor under the burden of many false myths.

As was mentioned before, railroads initially were subject to structural separation, but governments found it impossible to enforce this. (See Chapter 9 of [21] for more information about the transition of British railways to become carriers instead of just providers of rail roads.) And we have more modern evidence of the need for clear interfaces in telecommunications for effective regulation [11]. It appears very hard to regulate a packet data network like the Internet, which depends for its basic justification on statistical multiplexing. As has been pointed out before, the Internet has not been “net neutral” in the past, and it is not one now. The question of how one deals with legitimate network management issues, as well as with services that appear to reside most naturally in the network, is a thorny one. Extending (with appropriate modifications) traditional regulatory approaches to the Internet, as proposed, for example, in [41], may very well be the eventual outcome, but it won’t be easy, and will continue to provide lifetime employment to legions of lawyers and economists.

It would be far easier to enforce neutrality rules on a network that provided dark fiber or wavelengths only, which would be sold to service providers, who would then provide higher level services. That, of course, would require a complete restructuring of the industry. It might be wise, both for society and for shareholders of those companies, but very hard to achieve.

While technology imposes constraints on how effective regulation can be, it also imposes constraints on what service providers can do. Especially when the most valued services, such as voice, texting, and even simple search, are low bandwidth, it is hard for service providers to effectively regulate traffic (and thereby extract value), at least without very intrusive measures, which would likely be resented and could lead to protests and government intervention. The strategy of controlling just streaming video traffic is very unlikely to be productive.

9 Financial markets and the future of net neutrality

The importance of price discrimination in enabling deployment of infrastructure industries has been recognized by policy makers for a long time. Not only was it enshrined in the toll schedules of railroads, canals, and turnpikes, but there are historical examples where an increased level of discrimination was allowed in order to save a floundering enterprise. The paper [34] cites the Beverley Beck Navigation in the 18th century, in which the initial (already discriminatory) toll schedule was modified by the British Parliament to allow for a greater differential in charges for different goods. This was not an isolated case. In 1845, in order to fortify canals in their losing struggle with the rapidly expanding railway industry, Parliament passed a law explicitly authorizing canals to be their own carriers (something only a few had been permitted to be in the past), and giving them some leeway to vary their tolls (see Chapter 9 of [21]). (This attempt to keep the canals as viable competitors to railways did not succeed.) And as the Suez Canal was opening, when its financial prospects looked bleak, the *Economist* of Nov. 20, 1869 suggested that Suez tolls should be changed from the simple fixed fee per ton of cargo ship capacity and per passenger to ones that depended on the value of cargo. (This advice was not followed.)

Thus there are precedents for telecom companies to ask for ability to charge special fees to companies like Google that might be deriving large profits from the use of the infrastructure. The question is, do they need it? And there is little evidence that they do.

Extensive price discrimination was tolerated on railroads, even under regulation. However, railroads were a giant industry faced with demands for further expansion. In the U.S. in 1907 their revenues were almost 8% of the Gross National Product [9]. (By comparison, today telecommunications in the U.S. is at about 3-4%, depending on what is counted. Back in 1907, it was around 1.2%, if we count postal services, telegraph, and telephone, and the Post Office accounted for more than half of that total [29].) Furthermore, railroads were struggling with rising costs, as the demand for local transport grew, and rail technology was not suited to satisfy it. Telecommunications today is in a different environment, with most costs in a steep decline. This is only partly due to the Internet. Already almost a decade ago it was estimated that the entire plant of the U.S. phone companies, which had cost about \$340 billion, could be duplicated for about \$180 billion, or roughly half [37]. Further technological advances have taken place since then, so current costs would likely be lower, even in nominal dollars. (This evolution has been partially hidden by the rise of spending on wireless, which has kept total costs, revenues, and profits high.) Basically, of the three main segments of traditional telephony, access, switching, and long distance, which ages ago were regarded as about equally expensive, the second and third

have shrunk dramatically, and the access piece dominates. (For a discussion of how inexpensive the backbones of the Internet are, see [33]. And for evidence that those infamous “exafloods” are not swamping the Internet, see [7, 24].) That is why capital expenditure (capex) of the telcos has been weak, as their replacement cost has been consistently lower than their depreciation charges. And that is what has helped produce the bountiful free cash flow and profits of recent years.

Further evidence that telcos do not need vast new revenues to be able to afford a buildout of broadband networks comes from international comparisons. In many other countries far faster networks are available at lower prices. Even if one makes allowances for differentials in labor costs, taxes, and the like, it seems that broadband networks are all that expensive to build.

There certainly are some widely publicized claims that huge increases in investment are needed to provide widespread broadband connectivity. Probably the most extreme one is a study by R. N. Clarke of AT&T [8], which claims that what it calls “fully unmanaged IP networks” would require “between \$143 and \$416 per month” per connection. The detailed cost estimates in the Appendix to [8] are questionable. (As Dave Burstein has pointed out, if they were correct, AT&T’s U-verse service would be ruinously expensive to provide, contradicting what AT&T, Clarke’s employer, has been telling its investors.) But even if they were correct, they would be almost irrelevant. The basic assumption of [8] is that the “fully unmanaged IP networks” would need to be built today, with today’s technology and at today’s prices, and would need to accommodate essentially all foreseeable demands today. In particular, they would need to provide for transmission of all the video consumed today, but at higher resolution, all over the basic Internet access link. But in practice, Internet usage is not changing all that rapidly. Recent studies from Japan, which has bandwidths to the home comparable to what Clarke’s “fully unmanaged IP networks” envisage, show very slow spread of video [7]. The abstract of that study concludes that “observed trends ... suggest that video content is unlikely to disastrously overflow the Internet, at least not anytime soon.” And the title of that study talks of “slow crustal movement in residential user traffic.” Internet traffic growth is still fast, but has been slowing down, from about 100% per year towards about 50% in 2008 [24].

The economics of building new communications networks (or upgrading old ones, improved technology is allowing us to push far more data through the old twisted pair or coax connections that are in place) is very complicated, and this is not the place to get into it in detail. But there is far too little informed discussion of this topic in connection with net neutrality. Yet there evidence that costs of providing broadband connectivity, at current growth rates in data traffic, and current rates of technological improvement, are not huge. As just one example one can cite the projections made by Carphone Warehouse, an independent facilities-based ISP in the United Kingdom [6].

In general, projections by electronic and photonic equipment suppliers all suggest that current traffic growth rates of about 50% per year [24] can be maintained without dramatic boosts in capital investment.

Yet more circumstantial evidence that non-discriminatory communications systems should be viable comes from the wireline voice network. That is still the big revenue

producer on the wireline side, but operates in an exemplary net neutral fashion, and is, to an increasing extent, paid for by flat fees, those same flat fees that the telecom industry has disdained and misunderstood for over a century.

Some more evidence comes from the costs of the electric distribution system and others. In particular, the electric system reaches out further than the wireline phone system, yet is operated at a fraction of the cost of the latter (and in a very neutral fashion). Now of course electricity distribution does not have the switching costs of the telecom network, but it is those switching costs that are plummeting, as Moore's laws continue operating. What remains (relatively) expensive is just the cost of getting the physical conduit to the consumer.

The current campaign to justify high revenues and differential pricing on the Internet is reminiscent of the campaign in the late 1990s by telcos to charge special fees on dial Internet access. The argument then was that the long holding times on Internet access calls made with modems were wreaking havoc with the carefully planned telco equipment, and were making flat rate pricing unviable. Had regulators in the U.S. listened to such pleadings, the growth of the Internet would undoubtedly been greatly impeded. However, they instead turned a deaf ear and allowed flat rate to continue to apply for modem calls to the Internet. The result was a boom in profits for the telcos. In the first place, their complaints about increased traffic were very suspect, since that traffic was achieved largely through customers leasing extra lines just for Internet access, as is seen in Table 1 on p. 36 of [27] (although it should be mentioned that later in the 1990s, beyond the 1996 figure that was the most current at the time of writing of that report, average total length of calls per line did go up somewhat). Further, Internet access was concentrated later in the day (or night, to be precise) than voice calls, so did not interfere as much as claimed with basic traffic. And, finally, costs of switching equipment declined rapidly during that period. And so none of the predicted disasters materialized, and the industry thrived. Thus to correctly evaluate claims about need for additional revenues for broadband networks, one needs solid cost data and a dynamic model of the industry. At the moment we do not have either one available.

We should also remember that even with net neutrality rules in place, service providers would still be able to segment the market by speed of connections. Since the main value of data networks resides in provision of low transaction latency [28], not of capacity for streaming video at modest rates, this might very well be sufficient to differentiate among users, and yet still preserve the flat rate pricing those customers insist on.

While the telcos probably do not need to engage in differential pricing to pay for broadband networks, they may very well need such pricing to support their stock prices. But that is where we get into the great puzzle of the current financial markets. Returns on invested capital are very high, generally over 15% per year, a historical record. And they have been that high for several years, and they have been that high worldwide. This has been explained as a result of the entry of many developing countries, especially China and India, into the world economy. This essentially doubled the world labor force, while the stock of capital goods increased only a little. In this view, high equity returns are simply the outcome of a changed capital to labor ratio, with more going to pay for scarce

capital than for abundant labor. This is a reasonable argument, but it does not explain why long-term interest rates were as low as they were for a long time (on the order of 5% for high grade corporate bonds through early 2008. The contrast between 5 and 15% is striking. After all, equity and bonds are just two forms of finance, and while equity has traditionally earned more than bonds, it is hard to explain a gap this wide. Why doesn't someone borrow at the low interest rates and compete away those abnormally high equity returns? (The big borrowing binge by hedge funds and private equity funds was largely for financial investments, not for new brick and mortar facilities, nor for software development.) That was the big puzzle. The current (late 2008 and early 2009) financial crisis may be eliminate this puzzle. Another anomaly of the financial markets of the last few years was the unusually high fraction of corporate profits that were coming from the financial sector. But the huge writedowns by financial institutions in late 2007 and 2008, including several bankruptcies and wide government interventions, show that to a large extent those profits were a delusion, caused by defective accounting for value and risk.

One possible outcome of the financial crash might paradoxically be that it will encourage greater investment in telecommunications infrastructure. Even aside from government funding for economic stimulus, the crash might, after main turbulence subsides, lead to more realistic expectations of investment returns, which will make long-term investments in projects such as fiber to the home more attractive.

However the great puzzle of the financial markets is resolved, cable companies and telcos definitely contributed to it. As is described in [33], their valuations before the crash of 2008 were considerably above their replacement cost. (This was also true for most of other industries, result of those abnormally high equity returns.) As a rough approximation, it appeared that the valuations of both the cable and the telco sides assumed that they would get the entire telecommunications spending of the economy. That might be true for one or the other, but not for both, at least not unless they could somehow induce huge new increases in those expenditures, possibly by exploiting the potential of discriminatory pricing.

Policy makers in the U.S. have limited tools with which to induce deployment of broadband networks. But it is very doubtful whether imposing some sort of net neutrality would impede that policy goal, as the current rate of deployment seems gated by other considerations. Some form of net neutrality mandate, hard as it would be to define and enforce, might be more appropriate, in order to promote the development of the applications that will provide value to users and make them willing to spend on "dumb pipe" broadband, which might be much more effective in stimulating investment by current players. Left to themselves, current operators would surely concentrate on building out systems optimized for content delivery. Yet, as has been easy to predict from historical precedents [29, 30], the most promising avenues for stimulating interest in broadband by users is by promoting social interactivity. In that sense, success of services like YouTube and Facebook was very natural and predictable, yet it caught the industry by surprise. Thus simply providing more funding for current operators is likely to be wasteful, in that it would either be pocketed as extra profit, or spent in wasteful ways. The one thing that has been well documented (see, for example, [35]) is that established service providers are terrible at innovation in

services. Their core expertise is in widespread delivery of basic connectivity, and they, and their suppliers, have done well in innovating there, introducing DSL, cable modems, wireless transmission technologies, DWDM, and so on. But they have failed utterly in end-user services. The great success story to date in wireless data is texting, which was an accidental byproduct of the GSM technology development, not meant for consumer use. Even the one shining example of a successful content service, ring tone downloads, came as a surprise. At the 2004 OFC conference, where I pointed this out in a presentation, one of the later speakers, a prominent U.S. telco executive, was clearly offended, and responded with a declaration that the industry was very innovative. But the only example he could cite was Caller ID!

A structural separation, or strict net neutrality regime, could be beneficial for shareholders of the service provider companies, as it would enable those enterprises to concentrate on their core expertise, which is to provide basic connectivity, those much derided “dumb pipes.” Commodity services do not have to be unprofitable, after all. And recent moves, in particular the 2007 iPhone deal that AT&T concluded with Apple, may reflect the growing realization that the main sources of innovation in services are and will continue to come from outside. And with greater innovation, one can expect greater growth in demand for the basic transmission capacity (just as flat rate Internet access stimulated demand for basic voice lines in the late 1990s, and led to increased profits for telcos, even though they did not contribute a whit to anything on the Internet itself). But the vertical integration case will surely continue to be pursued because of its revenue and profit attractions. And note that the “cloud computing” idea is actually an extreme form of vertical integration, just carried out by other companies than the telecom service providers, and at higher levels of the protocol stack. There are attractions to end-users in some levels of vertical integration. The concept of a do-it-yourself end-to-end network is attractive, but few users have the skills and patience to make it a reality. So we will likely see extensive vertical integration, the only question is where, and what dangers will it produce.

The basic conclusion is that for pervasive infrastructure services that are crucial for the functioning of society, rules about allowable degrees of discrimination have traditionally applied, and are likely to be demanded for the Internet in the future. Those rules have often been set by governments, and are likely to be set by them in the future as well. For telecommunications, given current trends in demand and in rate and sources of innovation, it appears to be better for society not to tilt towards the operators, and instead to stimulate innovation on the network by others by enforcing net neutrality. But this would likely open the way for other players, such as Google, that emerge from that open and competitive arena as big winners, to become choke points. So it would be wise to prepare to monitor what happens, and be ready to intervene by imposing neutrality rules on them when necessary.

10 Acknowledgments

This research was supported partially by National Science Foundation FIND program grant CNS-0721510. Any opinions, findings, and conclusions or recommendations expressed in

this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

The author thanks the editors, the anonymous referees, as well as Dave Burstein, Richard Clarke, Fernando Gil De Bernabe, Stefano Quintarelli, and Hal Varian for their comments.

References

1. J. Alleman and P. Rappoport, “The future of communications in next generation networks,” white paper for the 2007 ITU Workshop on The Future of Voice, available at <http://www.itu.int/osg/spu/ni/voice/papers/FoV-Alleman-Rappoport-Final.pdf>.
2. J. Ankeny, “Survey: Texting remains top consumer draw,” *Fierce Mobile Content*, July 2, 2008, <http://www.fiercemobilecontent.com/story/survey-texting-remains-top-consumer-draw/2008-07-02>.
3. Anonymous, “The policy and prospects of the South-Eastern Railway Company,” *Monthly Railway Record*, March 1847, pp. 11–25.
4. J. Blau, “Threat of regulation prompts lower EU roaming fees,” *InfoWorld*, April 11, 2007. Available at http://www.infoworld.com/article/07/04/11/HNlowereuroamingfees_1.html.
5. S. Brin and L. Page, “The anatomy of a large-scale hypertextual Web search engine,” presented at the 7th International World-Wide Web Conference, Brisbane, Australia, April 1998. Full version available at <http://www7.scu.edu.au/1921/com1921.htm>.
6. Carphone Warehouse, Annual Strategy Day presentation, April 15, 2008. Available at http://library.corporate-ir.net/library/12/123/123964/items/287985/Analyst_Day_150408_Pres.pdf.
7. K. Cho, K. Fukuda, H. Esaki, and A. Kato, “Observing slow crustal movement in residential user traffic,” ACM CoNEXT 2008 conference, Dec. 2008. Available at <http://www.iijlab.net/~kjc/papers/kjc-conext2008.pdf>.
8. R. N. Clarke, “Costs of neutral/unmanaged IP networks,” white paper, Aug. 28, 2007 version, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=903433. Revised version in this issue.
9. U.S. Department of Commerce, Bureau of the Census, *Historical Statistics of the United States: Colonial Times to 1970*, 1975.
10. N. Economides and J. Tåg, “Net neutrality on the Internet: A two-sided market analysis,” available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1019121.
11. G. Faulhaber, “Policy-induced competition: The telecommunications experiments,” *Information Economics & Policy*, vol. 15, 2003, pp. 73–97. Available at <http://rider.wharton.upenn.edu/~faulhabe/Policy-Induced%20Competition.pdf>.
12. Federal Communications Commission, order for Comcast “to end discriminatory network management practices,” Aug. 1, 2008. Available at <http://www.fcc.gov/comcast-nr-080108.pdf>.
13. P. C. Fishburn, A. M. Odlyzko, and R. C. Siders, “Fixed fee versus unit pricing for information goods: competition, equilibria, and price wars,” *First Monday*, vol. 2, no. 7, July 1997, http://firstmonday.org/issues/issue2_7/odlyzko/index.html.

14. D. Gabel, *The Evolution of a Market: The Emergence of Regulation in the Telephone Industry of Wisconsin, 1893–1917*, Ph.D. Dissertation, University of Wisconsin–Madison, 1987.
15. B. Grow and K. Epstein, “The poverty business: Inside U.S. companies’ audacious drive to extract more profits from the nation’s working poor,” *Business Week*, May 21, 2007.
16. A. T. Hadley, *Railroad Transportation: Its History and its Laws*, G. P. Putnam’s Sons, 1885.
17. K. T. Healy, *The Economics of Transportation in America*, Ronald Press, 1940.
18. C. Hogendorn, “Broadband Internet: Net neutrality versus open access,” *International Economics and Economic Policy*, vol. 4, no. 2, 2007, pp. 185–208. Available at <http://www.springerlink.com/content/f2t5406412060407/?p=0e601bda14854b94b4465510b75ac69b&pi=4>.
19. ITU Workshop on The Future of Voice, 2007. Presentations and background papers available at <http://www.itu.int/spu/voice>.
20. M. W. Kirby, *The Origins of Railway Enterprise: The Stockton and Darlington Railway, 1821–1863*, Cambridge Univ. Press, 1993.
21. W. T. Jackman, *The Development of Transportation in Modern England*, Cambridge Univ. Press, 1916.
22. D. Lardner, “Railways at home and abroad,” *Edinburgh Review*, vol. 84, no. 170, Oct. 1846, pp. 479–531. Available online at <http://www.dtc.umn.edu/~odlyzko/rrsources/>.
23. D. Levinson and A. M. Odlyzko, “Too expensive to meter: The influence of transaction costs in transportation and communication,” *Phil. Trans. Royal Soc. A*, vol. 366, no. 1872, 2008, pp. 2033–2046. Available at <http://www.dtc.umn.edu/~odlyzko/doc/metering-expensive.pdf>.
24. Minnesota Internet Traffic Studies (MINTS) web site, <http://www.dtc.umn.edu/mints>.
25. A. Mohammed, “Verizon executive calls for end to Google’s ‘free lunch,’” *Washington Post*, Feb. 7, 2006. Available at <http://www.washingtonpost.com/wp-dyn/content/article/2006/02/06/AR2006020601624.html>.
26. A. M. Odlyzko, “The bumpy road of electronic commerce.” in H. Maurer, ed., *WebNet 96 - World Conf. Web Soc. Proc.*, AACE, 1996, 378–389. Available at <http://www.dtc.umn.edu/~odlyzko/doc/bumpy.road.pdf>.
27. A. M. Odlyzko, “The economics of the Internet: Utility, utilization, pricing, and Quality of Service,” unpublished manuscript, July 7, 1998, available at <http://www.dtc.umn.edu/~odlyzko/doc/internet.economics.pdf>.
28. A. M. Odlyzko, “The current state and likely evolution of the Internet,” in *Proc. Globecom’99*, pp. 1869–1875, IEEE, 1999. Available at <http://www.dtc.umn.edu/~odlyzko/doc/globecom99.pdf>.
29. A. M. Odlyzko, “The history of communications and its implications for the Internet,” 2000 unpublished manuscript, available at <http://www.dtc.umn.edu/~odlyzko/doc/history.communications0.pdf>.

30. A. M. Odlyzko, "Content is not king," *First Monday*, **6**, no. 2, February 2001, http://firstmonday.org/issues/issue6_2/odlyzko/.
31. A.M. Odlyzko, "Internet pricing and the history of communications," *Computer Networks*, vol. 36, 2001, pp. 493–517. Available at <http://www.dtc.umn.edu/~odlyzko/doc/history.communications1b.pdf>.
32. A. M. Odlyzko, "Privacy, economics, and price discrimination on the Internet," *ICEC2003: Fifth International Conference on Electronic Commerce*, N. Sadeh, ed., ACM Press, 2003, pp. 355–366. Available at <http://www.dtc.umn.edu/~odlyzko/doc/privacy.economics.pdf>.
33. A. M. Odlyzko, "The many paradoxes of broadband," *First Monday*, vol. 8, no. 9, September 2003, <http://firstmonday.org/issues/issue8.9/odlyzko/>.
34. A. M. Odlyzko, "The evolution of price discrimination in transportation and its implications for the Internet," *Review of Network Economics*, vol. 3, no. 3, 2004. http://www.rnejournal.com/artman2/uploads/1/odlyzko_RNE_sept.2004.pdf.
35. A. M. Odlyzko, "Telecom dogmas and spectrum allocations," written for the *Wireless Unleashed* blog, 2004, <http://www.wirelessunleashed.com/>. Also available at <http://www.dtc.umn.edu/~odlyzko/doc/telecom.dogmas.spectrum.pdf>.
36. H. J. Singer, "Net neutrality: A radical form of non-discrimination," *Regulation*, vol. 30, no. 2, 2007, pp. 36–41. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1001480.
37. B. Stuck and M. Weingarten, "A tale of two decisions," *Business Communications Review*, July 2002, pp. 47–48. Available at <http://signallake.com/publications/>.
38. B. van Schewick, "Towards an economic framework for network neutrality regulation," *Journal on Telecommunications and High Technology Law*, vol. 5, 2007, pp. 329–391. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=812991.
39. C.-C. Tschang, "China's Baidu.com accused of rigging search," *Business Week*, Jan. 12, 2009. Available at http://www.businessweek.com/magazine/content/09_02/b4115021710265.htm.
40. B. Warner, "Are the mobile networks backing the wrong horse?," *Times Online*, Oct. 17, 2007, http://technology.timesonline.co.uk/tol/news/tech_and_web/personal_tech/article2680101.ece.
41. P. Weiser, "The next frontier for network neutrality," *Administrative Law Review*, vol. 60, no. 2, 2008, pp. 1–50. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1080672.
42. E. Whitacre, "At SBC, It's All About 'Scale and Scope'," interview with *Business Week*, available at *Business Week Online*, Nov. 7, 2005, http://www.businessweek.com/magazine/content/05_45/b3958092.htm.
43. J. R. Wilke and D. Clark, "Microsoft is facing more telling e-mails in Minnesota lawsuit," *Wall Street Journal*, March 17, 2004. Linked online lawsuit evidence.
44. C. Wilson, "TELECOMNEXT: Net neutrality a bogus debate," *Telephony Online*, March 22, 2006. Available at http://telephonyonline.com/home/news/net_neutrality_debate_032206/.
45. C. Yoo, "Network neutrality, consumers, and innovation," Oct. 3, 2008 version, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1262845.