Economic Issues Facing the Internet

Hal R. Varian

June 10, 1996 (Revised September 15, 1996)

Contents

1 Introduction

The peculiar economic situation of the Internet is due to the fact that it existed for several years as a state-financed public utility. Because of this, the economic and business practices that would have normally been developed naturally as an industry grows and matures are missing from the Internet. The lack of these economic and business practices is by now painfully clear and will become more so in the near future.

However, the flip side of this issue is that the Internet would probably never have developed as a purely private industry. The open standards and the public-spirited pattern of cooperation that have characterized the Internet's history would probably not have been possible in a privatized environment.

Now these ''public'' and ''private'' facets of the Internet are increasingly coming into conflict. The key problem facing us will be how to maintain the openness of the public model while, at the same time, developing the sustainable business models necessary for future growth and development.

2 The past: a network for scientific communication

To understand the present and future of the Internet, one must first understand its past.

The Internet protocols were developed by the Advanced Research Projects Administration (ARPA) as part of an effort to design a robust communications network. The ARPAnet, which linked together a number of high tech research institutions, was deployed both to demonstrate the workability of the protocols and to facilitate communication among research communities.

In the early eighties, the National Science Foundation (NSF) created several supercomputer centers around the United States. Researchers at universities that did not have local access to a supercomputer wanted remote access. To provide such access the NSF funded the deployment of a network based on Internet protocols whose purpose was to connect all major research universities to the supercomputer sites and to connect the supercomputer sites to each other.

The NSFNET backbone consisted of several leased telephone lines that connected the supercomputer centers to each other. The NSF also provided funds to subsidize the development of regional networks that connected research universities and institutions to the NSFNET backbone. The NSF directly financed most of the costs of the NSFNET backbone; the costs of the regional networks were covered by NSF subsidies and contributions from the research institutions connected to them. Furthermore, colleges and universities could apply for grants to the NSF that helped cover some of the costs of connecting to their regional network.

Although the initial purpose of the NSFNET was to carry out research at the supercomputer centers, it soon became apparent that many more uses of the network were possible. A fundamental property of networks is that if we are all connected to the same network, then we are all connected to each other. This meant that email, file transfer, remote login, and other tools could be used to communicate among universities about matters having nothing to do with the supercomputer centers. By the late eighties, the Internet was a rousing success due to applications that were totally divorced from its initial rationale.

2.1 Administration of the NSFNET

For the first part of its life the NSFNET backbone was run by Merit[1] (Michigan Educational Research and Industrial Triad), a nonprofit regional network for Michigan. The NSF award[2] to Merit was based on its standard evaluation policies for scientific projects. During this period, the NSF promulgated an Acceptable Use Policy[3] which restricted the use of the Internet to non-profit activities.

In the late 80s, several for-profit concerns wanted access to the Internet so that they could communicate with users, researchers and customers at universities. The NSF did not want to be in the business of subsidizing Internet access by for-profit firms, so they instituted a change in the administrative structure of the NSFNET. Merit, IBM and MCI each contributed resources to a new non-profit entity called Advanced Network Systems (ANS). Merit contributed technical know-how, IBM contributed workstations and hardware, and MCI contributed networking equipment and discounts on telecommunications lines. Merit then subcontracted with ANS which, in turn, managed the day-to-day operation of the Internet. Merit also provided a number of higher level services such as documentation, user education and marketing, record keeping and research efforts on networking.

ANS was then free to sell backbone access to all potential users, be they nonprofit or for-profit entities. Meanwhile, the demand for commercial access to Internet access led to the formation of a few for-profit backbone networks such as UUNET and PSI (Performance Systems International). These companies sold services primarily to commercial concerns but interconnected with both the regionals and the NSFNET backbone at various points. The network topology allowed all Internet users in the US to interconnect, regardless of whether they were government, university, or private concerns. (See Merit's NSFNET Backbone Project[4] for a more detailed account of this history.)

The NSF also ran an international connections program that encouraged connections with other countries' scientific research networks. A typical arrangement was to split the costs of connection equally. In some cases, a country's connection to the Internet was a by-product of US scientific projects in that country. For example, the US Antarctica expeditions were headquartered in New Zealand. When the Antarctica teams wanted Internet connectivity they were able to share the costs with the New Zealand research establishment.

[1] http://www.merit.edu
[2] http://www.nsf.gov/ftp/awards87/awd8720/a8720904.txt
[3] http://www.merit.edu/nsfnet/acceptable.use.policy
[4] http://www.merit.edu/nsfnet/.index.html

2.2 Privatization

By the mid nineties it became clear that there were many independent companies who were interested in providing Internet services. In particular, it was clear that what was once an experimental technology was now a commercially viable business. The NSF had built the Internet, but now that it was in existence, there was little reason for the NSF to continue to pay for its upkeep.

In 1993 the NSF issued a ‘‘request for proposals’’ that outlined the structure of a new, privatized and commercialized Internet. The privatization plan[5] had four main components:

- The NSF would eliminate all support for the NSFNET backbone.
- The NSF would phase out financial support for the regionals.
- The NSF would help fund the creation of NAPs (Network Access Points) at which multiple backbone providers could interconnect.
- The NSF would fund development of a very-high speed network (vBNS) which would serve both as a testbed for network research and provide support for scientific applications.

There was a certain amount of anguish among academic circles about ‘‘shutting down the Internet.’’ However, these protests were mostly due to misinformation. Even in its heyday the NSF spent only $12 million or so on the backbone and $8 million on subsidies for the regionals and related programs. If there were, say, 10 million potential users at this time, the NSF subsidy only amounted to $2 per user. When the NSF stopped subsidizing the backbone, costs for Internal access for a large university about doubled from around $60,000 per year to $120,000 year. But in a university with 30,000 students, this is only an extra $2 per student. This expense was dwarfed by the costs of internal networking and support, and posed little additional burden on the scientific and research community.

3 The present: a network for the digerati

The NSFNET backbone was shut down on April 30, 1995, and the transition to the new privatized network went relatively smoothly. MCI and SPRINT made major efforts to sign up new customers for their backbone services; it now appears that MCI is the largest carrier of Internet traffic in the US. There have been some new entrants into the backbone provision market, most notably AT&T.

There were 3 new NAPs established by the NSF program; one in Chicago (run by Ameritech), one in San Francisco (run by PacBell), and one in New York/New Jersey (run by AT&T). Along with MAE-East (Washington) and MAE-West (San Jose), run by MFS Datanet, these served as default interconnection points. Large backbone providers also established other interconnection points at convenient locations, but

[5] http://www.merit.edu/nsf.architecture/.index.html

their locations are generally not part of the public record. See Merit's Overview of the New Network Architecture[6] for more details.

Some of the NAPs chose tried-and-true networking technology; others were more adventurous, choosing cutting-edge technologies such as ATM. As it turned out the cutting-edge choices did not work well and most NAP operators ended up falling back on conventional technology such as FDDI. The NAPs have ended up being significant congestion points for the Internet since so much traffic is interchanged at these locations.

The regionals were originally closely connected to universities: BARRNET (the Bay Area regional research network) was housed at Stanford; NEARNET (the New England regional) was at MIT, and so on. Several of these regionals were sold to a private concern (BBN), while others remained under non-profit control (e.g., MichNet, the Michigan regional).

The transition to the new administrative structure took place at the same time the Internet was experiencing accelerating growth. Although national statistics are no longer kept, it is widely thought that Internet traffic has *more than doubled* in the last year, reflecting an acceleration in the rate of growth. This dramatic increase in traffic is due to a great extent to the introduction of the World Wide Web, a new application that has achieved widespread popularity.

3.1 Network infrastructure

3.1.1 The World Wide Web

The WWW was developed by Tim Berners-Lee, a computer programmer at CERN, a research lab for high-energy physics in Switzerland. Berners-Lee's idea was to use hypertext to link together different documents on the Internet. The first version of the system he released had a graphical user interface developed on the NeXT computer. This was quickly followed by an ASCII interface that could run on any Unix computer.

The World Wide Web might have remained a system used only by high-energy physics researchers, were it not for the efforts of a small research team at the National Center for Supercomputer Applications. The NCSA, at the University of Illinois, was one of the supercomputer sites that the NSFNET was originally set up to support. They developed a number of applications to assist researchers who used the Internet, and in 1993 they released the first version of Mosaic, a tool that provided a unified interface to several different protocols used on the Internet.

Mosaic added an important feature to the Berners-Lee conception of the World Wide Web: images. The graphics support was relatively primitive by today's standards, but it added a whole new dimension to exploring the Internet. It is insufficiently appreciated how important the addition of images was to the spread of the Web. Pictures added the ''gee whiz'' factor that made people excited about the Web. Subsequently, when the Internet began attracting significant commercial interest, companies were able to use corporate logos and other methods of maintaining brand identity. Without the ability to display images, it is unlikely that the WWW would have achieved the kind of presence it now has.

[6]http://www.merit.edu/nsf.architecture/.index.html

The downside of the success of the the WWW is that pictures take up a lot of bandwidth. All those corporate logos have a cost in terms of bits, which is one reason why traffic on the Internet has increased so dramatically.

3.1.2 Multiple service providers

The Internet has been suffering growing pains, and for good reason: the administrative structure underwent a significant transition at the same time that the way the Internet was used significantly changed. One of the significant problems facing today's Internet is that many of administrative and business procedures that were designed in the days of a single, non-profit backbone are no longer appropriate for the new multiple, for-profit backbones.

If you have multiple backbones, it is generally necessary to exchange traffic across several networks. For example, sending a packet of data across the country can easily involved 15 or 16 hops that traverse half-a-dozen independently owned networks.

This involves two sorts of costs. One cost is the cost of maintaining the ''routing tables.'' These are large tables of data that describe the paths that packets take as they traverse the Internet. Keeping them up-to-date and optimized is a non-trivial task. It is very tempting to let ''someone else do the job'' and simply pass traffic along to someone who is known to have a good (or at least adequate) routing table. Back in the old days, this was the NSFNET, which, being essentially a public utility, was willing to accept the responsibility of maintaining routing tables. However, nowadays it is widely thought that (at least some) networks are not investing adequate resources in maintaining good routing tables.

Another issue involves quality of service. A network provider can use faster or slower routers, configure larger or smaller buffers, etc. Many new applications such as WWW, Internet audio, CUSeeMe, etc. require significant bandwidth. But just as a chain is only as strong as its weakest link, the network is only as fast as its slowest segment. If a packet of information traverses several different networks, it is the slowest one that serves as a bottleneck.

Right now, there is no way for a packet to request ''priority'' or preferential treatment. All packets are treated the same. However, applications vary significantly in the demands that they place on the network: real-time audio or video places significantly greater demands than email. There is no way to indicate demands for different priorities now, and applications that require high-quality service are not very reliable.

The rapid growth in the World Wide Web has contributed to a significant congestion problem at several locations on the Internet. The trans-Atlantic and trans-Pacific links have been persistently slow, and the NAPs often become overloaded. Graphics-intensive Web pages look cool---but they take forever to load and add to congestion on the net.

3.2 Content

Up until now I have discussed issues involving the basic Internet infrastructure---the plumbing. I now turn to a brief survey of the current state of the content available on the Internet.

The largest amount of traffic on the Internet during the NSFNET era involved file transfer. The files were often text files consisting of computer programs, academic papers, and the like. During the last few years of the NSFNET era, more images were transferred which consumed significantly more bandwidth. (See Schwartz94[7] for an analysis of Internet traffic patterns.)

Today, it appears that Web traffic may be the largest form of Internet activity in terms of bits transferred. It is hard to be exactly certain about this since there are no longer publicly available statistics on the amount and composition of Internet traffic.

There is now significant interest in using the Web as a publishing medium. Many magazines, TV stations, academic journals, and commercial enterprises of all sorts now maintain a presence on the Internet.

Several forms of Internet payment systems are under development. See my Information Economy/Commerce page[8] for a list of over 20 such systems. Each touts its various advantages or disadvantages: anonymity, audit trails, reliability, compatibility, etc. There are several Internet payment systems now in use, but it appears likely that there will soon be a shakeout that will leave only few standard methods.

In addition to the development of payment mechanisms, there is also much experimentation with pricing policies. Will information be sold by subscription, by the ''piece'', or supported by advertising? It seems a safe guess that all of these will co-exist, so the real question should be *which* kinds of pricing are a appropriate for which kinds of information.

One of the most active areas of commercial interest has been in the area of information discovery. Indeed, services such as Yahoo[9], Lycos[10], and other Internet cataloging and classification services have attracted great commercial interest.

The big question that commercial interests face is that of asking what kind of information users will pay for, and how much they will pay. Some examples seem clear: user support such as that provided by Federal Express[11] is certainly attractive. Virtually any company that offers a menu-driven voice access to information would be well-advised to set up an Internet site for this form of customer support.

However, after that observation, the commercial potential of the Internet becomes much more murky. Will people really read newspapers and magazines on their computers after the novelty wears off?

3.3 Policy

With respect to policy, there are several issues looming on the horizon.

---

[7] ftp://ftp.cs.colorado.edu/pub/techreports/schwartz/FTP.Meas.ps.Z

[8] http://www.sims.berkeley.edu/resources/infoecon/Commerce.html

[9] http://www.yahoo.com

[10] http://www.lycos.com

[11] http://www.fedex.com

3.3.1 Decency

The bits hit the fan during the summer of 1995 when Congress discovered that one could access dirty pictures on the Internet. The resulting Communications Decency Act of 1995 was an attempt to remedy this problem. The CDA makes it illegal to use an interactive computer service to display to a person under 18 years of age any ‘‘comment, request, suggestion, proposal, image, or other communication’’ that is ‘‘patently offensive as measured by contemporary community standards’’. (See the text of the Communications Decency Act[12].)

It is widely believed that the CDA will not survive court scrutiny, which is, in this author’s opinion, appropriate. (See the Electronic Frontier Foundation’s homepage[13] for an analysis of the CDA[14].) However, one cannot ignore the fact that there *is* a problem. Some material on the Internet is not appropriate for children, while some is. How can a well-meaning parent (or schoolteacher, or librarian) deal with this issue?

A system called PICS[15] (Platform for Internet Content Selection) may provide the answer. PICS defines a set of protocols that allow the user to designate ‘‘rating agencies’’ that will provide ratings of different Internet sites. The user can then configure his browser, or his children’s browser, to display only content that is deemed appropriate. The PICS system itself doesn’t define who the rating services are, or what criteria they use. Anyone can provide PICS ratings. It also doesn’t define how the ratings will be used: that’s up to the browser provider. All that PICS does is describe a set of protocols that can be used to convey ratings to users.

This highly decentralized approach to content rating is very appealing to the Internet philosophy and many firms have joined the PICS project including Apple, America Online, DEC, IBM, MCI, Microsoft, Netscape, Prodigy, etc. As we will see below, PICS may end up being used for many other services than simple content rating.

3.3.2 Encryption policy

Another issue that is currently being hotly debated is encryption policy[16]. The issue here is that it is now technically feasible to encode documents in a way that is widely believed to be ‘‘unbreakable.’’ According to some, this would allow criminals and terrorists to encode their communications in a way that would prevent detection and prosecution. Accordingly, the US has banned the export of the algorithms that allow for this type of encryption.

This policy has had essentially no effect on the dissemination of information about encryption algorithms. The basic algorithm is so short and simple that it can be printed on a T-shirt (giving a new twist to the term ‘‘software’’), and it is virtually impossible to control the spread of this kind of information.

Regardless of the merits or demerits of controlling encryption technology, it seems that it is essentially impossible to do, and it is widely expected that the U.S. government will eventually give up on the effort.

[12] http://www.eff.org/pub/Alerts/s652_hr1555_96_draft_bill.excerpt
[13] http://www.eff.org/
[14] http://www.eff.org/pub/Alerts/cda_020296_eff.statement
[15] http://www.w3.org/pub/WWW/PICS/iacwc.htm
[16] http://www.sims.berkeley.edu/resources/infoecon/Intellectual_Property.html

3.3.3 Telecommunications deregulation

The huge growth in the Internet is coming at the same time as significant telecommunications deregulation. The Telecommunications Act of 1996 left many ambiguities which are supposed to be hammered out the the FCC in the next year. Several of these are of direct relevance to the Internet.

In Spring 1995 the America's Carriers Telecommunications Association (ACTA) filed a petition with the FCC which urged them to regulate voice on the Internet. (See ACTA Petition Resource Center[17] for details.) The Internet is currently classified as an Enhanced Service Provider by the FCC and is therefor not regulated. The ACTA petitition argued that ''Internet telephony'' is fast becoming a reality and that the (US portion) of the Internet be subject to regulation. In particular, Internet service providers should pay access charges to the local exchange carriers for use of their facilities. These access charges, paid by all long-distance carriers, are on the order of 5 cents a minute.

Regardless of whether one thinks that this is a good idea or a bad idea one is left with the fundamental question: how could such charges be measured? Only a small fraction of Internet users have experimentd with voice applications, There is essentially no easy way to tell if a given packet contains textual data, voice data, or video data---it's all just bits. This is especially true if the data is encrypted, as is likely to be the case in the near future. Even if one *wanted* to regulate voice on the Internet, it appears that there is no easy way to do it.

Another issue is capacity utilization of the local exchange. Many people access the Internet via POTS (plain old telephone service) and this ties up circuits in the local loop for much, much longer than typical voice calls. The Bell Atlantic Report on Internet Traffic[18] indicates that the average length of calls to an Internet Service Provider is around 18 minutes, as compared to the average call of 4-5 minutes for other calls. According to this report, if just 15% of the households in the calling area went on the Internet at the same time and stayed on for an hour, local network switching capacity would have to *double*! Satisfying this demand would involve significant investment in POTS capacity, which is not really a long-run solution for Internet access.

However, there is some hope that technology will bypass regulation. Providers of low-level satellite wireless communication systems are building Geographic Positioning Systems into their hand-held phones. The reason for incorporating GPS is that if the wireless provider knows the location of a call they can then set a price that is competitive with other wire-based service providers. If you call from the home or office---where you have access to standard telephone service---the rates will be low, to reflect the fact that competition is present. If you call from your car or boat, rates will reflect the monopoly conditions.

The bottom line, from the viewpoint of regulatory policy, is that the local providers may well face significant competition from wireless, especially for low-bandwidth applications.

---

[17] http://www.cais.net/cannon/acta.htm
[18] http://www.ba.com/ea/fcc/report.htm

3.3.4 Intellectual property

Congress is considering a set of revisions in the intellectual property law that may be quite significant. The proponents of these changes argue that they simply clarify existing law so that it applies to the new digital world. Detractors argue that the changes are a major departure of current intellectual property case law and question whether these changes are either necessary or appropriate.

There are several technological solutions that may help control intellectual property on the Web. One popular research scheme involves ‘‘software envelopes’’ or ‘‘cryptolopes’’ that contain encrypted versions of the material to be displayed. These envelopes are designed so that when they are ‘‘opened’’ the user is automatically billed for viewing their contents. (See my Web page on intellectual property[19] for other examples of intellectual property protection.)

Many people believe that this is the ultimate solution to intellectual property control, and support for cryptolope systems of this sort is built into the proposed revisions in the intellectual property law. (Material on the Lehman ‘‘White Paper’’ describing such revisions is available on my intellectual property Web page[20].)

I think that there are certainly applications where this sort of thing may be useful, but I have doubts about whether their use will really be widespread. In the mid-eighties there was much concern over illicit copying of software. Companies came up with all sorts of software protection schemes that, for the most part, inconvenienced legitimate users and had little effect on illicit users. Competition effectively eliminated most forms of IP protection that involved inconvenience. The current systems are simple digital watermarks (name and employer watermarked into each copy of the software), with some special protection for niche market products.

It is likely that this will be the outcome of the current debates as well. Anything that inconveniences the user will not be tolerated. Simple digital watermarks will suffice for most mass-market products. For special purpose products software envelopes and the like will be used, but this will be the exception rather than the rule.

Another important consideration is that many of the things that producers would *like* to do with with software envelopes and other exotic technologies may be easier to accomplish with simple pricing strategies. Just as most phone calls are within the organization, much sharing of intellectual property occurs within an organization. Site licenses, subscriptions, and other forms of group-based pricing can work well for these kinds of intellectual property.

Esther Dyson has suggested that the best way to develop a business model for distributing intellectual property on the net is to treat content *as if* it were free. This is not to suggest that content is, or should be free---rather, this thought experiment leads one to consider ways of adding value to content by utilitizing the unique characteristics on digital access: hypertext, cheap search, timely delivery, and so on.

[19] http://www.sims.berkeley.edu/resources/infoecon/Intellectual_Property.html
[20] http://www.sims.berkeley.edu/resources/infoecon/Intellectual_Property.html

4 The near future: a network for everyone

We turn now to the most dangerous part of the this essay: short-term forecasting. Describing the past and present is easy, and describing the distant future (5 years hence) is safe since no one will come back to check. But the near future---the next year or two---is problematic.

4.1 Network infrastructure

The first prediction is that we will see a shakeout in the backbone industry. This is due to the simple fact that the price at which a telco will lease a line is much, much greater than the cost to the telco of providing that line. Hence telecommunications companies (and other firms that own their own lines) have a significant cost advantage over those who have to lease them. This means that the MCIs, AT&Ts, and Sprints will end up dominating the ANSs, PSIs, and UUNETs. Indeed, several of these latter companies have already been acquired or are likely takeover targets. Backbone provision is likely to become (some would say, remain) a commodity business which is dominated by a few large firms.

Another possibility would be a price war in leased fiber and telecommunications services in general. Many industry observers view the AT&T-MCI-Sprint long-distance service industry as an oligopoly. If the oligopoly cracks, prices could fall dramatically, which would mean that the lease prices of line could end up approximating their costs. This might allow independent providers to survive.

However, even in this scenario there is a problem: quality of service. New Internet applications are demanding higher and higher qualities of service. Current protocols and practices are not up to the task. One possibility is for new protocols and practices to be developed that would allow users to specify a priority for packets. This would allow packets carrying voice and video to have higher priority than email.

It's easy to design protocols like this. The hard thing is getting them deployed. Deployment requires that everyone agree on how priorities will be denoted and handled. But getting agreement among a very disparate group of backbone, regional and ISPs will be very difficult. A more likely scenario is that the larger providers will begin to offer end-to-end performance ''guarantees''. This means that video conferencing between two AT&T customers can be of adequate quality but video conferencing between arbitrary customers will have inferior quality. The fact that larger companies can then provide better quality service will lead to an increase in their market share. Logically this could result in only one or two companies providing end-to-end service. More practically, antitrust concerns would result in 3-5 end-to-end providers in the US, with most other countries having only a single provider.

4.1.1 Cable modems

Cable modems offer the promise of high-speed (1-2 Mbs) access to the home. There are three big problems facing their deployment. First, no one knows if consumers are willing to pay the requisite amounts for such access speeds. Whether they will or not will depend on the content available and the ease of use of the resulting systems.

Second, no one really knows if cable modems will work. They work fine in small scale trials, but network technology is well-known to have problems in scaling. Many existing cable lines are rather noisy; this is tolerable for analog signals, but these lines will require upgrading to handle digital signals, which may be a significant additional expense. Furthermore, the current plans for cable modem deployment call for bandwidth to be shared among several users. If everyone watches the Superbowl updates at the same time, performance could be severely degraded.

Finally, there is the issue of backbone meltdown. If there is really significant demand for high bandwidth to the home, say for videoconferencing, the current backbones simply will not be able to handle it. Many companies seem to be designing their Internet strategy around forecasts of unlimited bandwidth. Although significant increases in bandwidth will likely become available eventually, dramatic increases are unlikely to be available in the near-term future.

Although everyone seems to be focussed on the home market, the business market is likely to be the biggest demander of bandwidth in the near future. The phenomenal popularity of email and Intranets is just the beginning. Video and audio conferencing may well become major applications in the near future.

In order to support this demand for bandwidth, companies will need to be able to contract with network providers who can give them adequate performance, which likely requires end-to-end connectivity. This is likely to be a significant business opportunity for network providers like AT&T, MCI and @home.

4.2 Content

In the next two years we will see secure payment mechanisms become commonplace on the Internet. The answer to ‘‘which ones’’ is simply whichever ones VISA and MasterCard endorse. This statement illustrates my belief that anonymous payment mechanisms (like Digicash) will probably remain niche players. See my Web page on Internet Commerce[21] for information about several sorts of anonymous and non-anonymous commercial protocols.

The reason is that the non-anonymous nature of credit card purchases is highly valuable due to the marketing data it provides. Companies can use information about spending patterns in a variety of ways, and they would therefore be willing to pay a premium for non-anonymous transactions. Consumers, on the other hand, *say* they value their privacy, but I suspect that they would be willing to part with it for rather small amounts. I thinks my privacy is important, but, hey, give me a few frequent flyer miles and I’ll tell you anything you want to know!

4.2.1 Pricing information

There is a lot of confusion these days about how to price information, Some think it should be paid for by the byte, some by advertising, some by subscription. The simple answer is that all of these pricing plans will be used---and maybe more.

[21] http://www.sims.berkeley.edu/resources/infoecon/Commerce.html

Many observers think that ‘‘by the byte’’ pricing is likely to be a significant form of pricing. I expect the contrary: subscription-based pricing is likely to be the most significant way that information is sold. This is due to a fundamental property of information: you don’t know if you want it until after you’ve seen it . . . and by then you’ve already paid for it. This means that there is a large ‘‘reputation effect’’ in information---I read the *New York Times* today because I liked what it said yesterday. The same things goes for most other forms of information I use: I tend to use the same sources since they have established a reputation (with me) for usefulness, reliability, and entertainment.

Now it is true that this feature does not in itself require subscription pricing. I purchase books by my favorite authors at bookstores, not via subscription. However, when information is delivered electronically, it offers the possibility of personalization. As I mentioned before, producers can offer a better product when they know to whom they are selling it. A repeated relationship---such as that offered by a subscription---allows for a more personalized product, benefiting the producer and consumer alike.

Hence subscriptions will sell at a discount---as they do today---and they will become a popular form of paying for information---as they are today.

When we put together the possibility of an ‘‘information relationship’’ with the online delivery, we are immediately led to ‘‘mass customization’’ of information goods. There will be several different forms of each information product and consumers can choose the one they prefer. Some of these will be pre-constructed, to appeal to different demographic groups, others will be personalized on the fly. Indeed, information and communications technology are the enabling technologies allow for mass customization of physical goods as well.

Pricing considerations play an interesting role in this product-line development. It is widely recognized that information goods involve high fixed costs and low incremental costs. Such a cost structure is naturally associated with price discrimination---charging different prices for the ‘‘same’’ product. ‘‘Same’’ is in quotes since the product is often modified in order to support the different prices.

Books are issued first in hardback, then in paperback, and finally may be remaindered. Stock market quotes cost one price if they are 20 seconds old and a much lower price if they are 20 minutes old. Online services, such as DIALOG, charge nearly every user a different price.

The strategy of changing the physical form of the packaging is not available for purely digital information, but delay is certainly a viable strategy. Other ways to differentiate information are with respect to ease-of-use, resolution, whether it is displayed on the screen or printed, etc. All these methods and more may be used to differentiate the underlying information good and support price discrimination.

4.2.2 Information retrieval

The fundamental constraint in the information economy is the limited information processing power of the human brain. The most difficult problems we face will be in finding useful, valuable, and relevant information. I expect to see some significant advances in this area.

First, there will be smarter systems with much friendlier interfaces. User obsequious interfaces will be the norm.

More importantly, I expect to see the emergence of a variety of Better Bit Bureaus---systems that provide recommending services. Some of these will be catalogs such as Yahoo[22], others will be search engines like Alta Vista[23].

Another very interesting form of recommendation systems are social recommenders[24]. These are electronic versions of systems like Zagat's restaurant guide or *Consumers Union* movie reviews that utilize the collective experiences of their participants to evaluate items. In GroupLens[25], for example, participants rate materials they read. When you are presented with a list of items to examine, you see a weighted average of the ratings of previous readers. So far, this is quite straightforward. But here's the gimmick: the weight that each person receives in the weighted average depend on how often you have agreed with that person in the past. Cute, isn't it?

In systems of this sort you can ''discover'' someone with common interests and judgements---without ever meeting them! This allows for the emergence of virtual communities with common interests, electronic panels of experts, peer refereeing, and so on at much lower transactions costs than current practices.

The PICS system mentioned earlier builds in enough infrastructure so that above means that anyone will be able to set themselves up as a recommender and transparently provide guidance to users. This could be specific groups (like *Consumers Union,* or your local library) or social recommenders like Zagat's or GroupLens.

Once systems of this sort are widely available, they can play the role of evaluating content for appropriateness and for interest. The nice thing is that rather than have one organization rating movies P, PG, R and X, you can choose the recommender of your choice---or any collection of them for that matter.

4.3 Policy

4.3.1 Intellectual property

I have already mentioned that I expect that intellectual property issues will fade. This is because consumers will reject any intrusive or difficult-to-use form of software envelopes and suppliers of information will come up with more effective ways to price information. Software envelopes will be used for some sorts of content, but watermarking, subscriptions, and site pricing will be the more common ways to deal with IP problems.

Of course, I am describing IP problems in the developed country context. The situation in China or India is a different matter. There, widespread counterfeiting is commonplace and IP protection will be very difficult. One of the fallouts of loose IP protection is that there is little incentive to produce software with local content. There are few Hindi wordprocessors, although there is a lot of programming talent in India. As countries become richer, they want to provide more of their own content---and to do so

[22] http://www.yahoo.com
[23] http://www.altavista.digital.com/
[24] http://www.sims.berkeley.edu/resources/collab/
[25] http://www.cs.umn.edu/Research/GroupLens/

requires some degree of adherence to intellectual property safeguards. The comparison with environmental issue is illuminating: if per capita income is below $4,000 counties choose the cheapest technology, which is often the dirtiest technology. Once per capita income rises above $4,000 countries become much more environmentally aware---for the simple reason that citizens want to live in a clean environment. The same thing is likely to apply with intellectual property. As developing countries become richer they will begin to take intellectual property seriously.

4.3.2 Privacy

The big issue of the next two years is likely to be the one of privacy. Just as ''decency'' dominated public discussion in 1995, privacy will be the big issue of 1997--98.

Recall that I said above that Americans value their privacy, but they may be willing to sell it for a low price. The policy issue that that needs to be confronted is defining a property right in privacy so that people have something to sell.

Privacy has many dimensions, but in this chapter we are only interested in the economic aspects which involves the relationships between buyers and sellers. The three important aspects of privacy that are relevant here are ''annoyance'', ''actuarial information'', and ''willingness to pay.''

We first consider the annoyance dimension. I don't care if you know my phone number, as long as you don't call me and try to sell me something. The same thing goes for my email address. This desire not to be annoyed is relatively easy to deal with since the incentives of the producers and consumers are aligned. The producer doesn't want to advertise to someone unlikely to buy its product, and the consumer may *want* to see information that will improve his or her purchasing decisions. Thus both sides of the market have an interest in seeing that relevant information is available and that the producer does not unnecessarily annoy the consumer. It may be necessary to have legislation about what kinds of advertising practices are appropriate and which aren't, but it should not be terribly difficult to achieve consensus on this matter.

The second category is ''actuarial information.'' For example, an insurer may want information about the health or the sexual practices of someone they are considering insuring. Similarly, a loan officer at a bank may want to know certain information about the individual. In this case, on party wants information that the other party doesn't want him to have, so incentives are perfectly *misaligned.* However, note that parties who have good health practices (e.g., nonsmokers) or are good credit risks would be quite willing to reveal this information, since it will distinguish them from the bad risks.

The third dimension of privacy relevant to commercial transactions is the desire of one party to know the reservation price (maximium willingness to pay) of the other party. More realistically, one party in a transactions would like to know *characteristics* of the other party that are known to be correlated with willingness to pay (e.g., demographic characteristics, possesions). Here again, interests are diametrically opposed.

In the cases where interests are diametrically opposed it appears that the government must step in to define the proper line between the public and private spheres. Defining this line will be very complex and contentious and we can expect that it will take some time before firm groundrules are in place.

5 The year 2000: the network matures?

It is time for some speculation about the far distant future: say 4 or 5 years from now. It is said that ‘‘Internet years’’ are like ‘‘dog years’’----they move 7 times as rapidly as human years. So forecasting what will happen to the Internet in 4 or 5 years is like forecasting 30 years in the future . . . and is likely to be just as accurate.

5.1 Network infrastructure

I suggested earlier that there would be a shakeout in backbone providers and concentration in a few large firms. By the year 2000 I would expect to see significant concentration with a few large firms offering end-to-end service along with reasonably well coordinated practices for exchanging traffic among themselves. Settlement schemes that support the quality assurance necessary to support real-time applications will be worked out among these large organizations. These contractual arrangements will effectively preclude new entrants.

A side benefit of providing end-to-end service is that we will have the equivalent of ‘‘Internet dialtone.’’ The hassles of connecting to the Internet will be a thing of the past, and it all will ‘‘just work’’ as a matter of course.

5.2 Content

5.2.1 Browsers and operating environments

By the year 2000 browsers such as Netscape will either disappear or will be incorporated into the applications. Here there are two scenarios. One I call ‘‘Netscape becomes the operating environment’’, the other is ‘‘Gates controls the doors’’. In the Netscape scenario, the present trend towards add-ins and helper applications continues and Netscape develops into an entire operating environment for applications. The browser essentially becomes a network file browser and applications operate within this environment.

The other scenario---which I consider more likely---is that Microsoft incorporates Internet awareness into all of its applications, right down to the operating system itself. Since all applications are Internet-aware, there is no need for an independent browser.

The reason why the ‘‘Gates controls the doors’’ is the more likely outcome is that users are wedded to their applications. It is much easier to extend the functionality of an existing application than it is to get users to adopt yet another set of applications and user interfaces.

5.2.2 Advances in ergonomics

The next big surge in content provision on the Internet will come when computers become significantly easier to manipulate. Right now no one can read material on a computer screen for enjoyment---the ergonomics are too unpleasant. When 300 dpi screens become available that are the size and shape of a book, all this will change. Such ‘‘dynabooks’’ will lead to significantly improved ergonomics. The user can sit however they want and move the dynabook to a comfortable location. They no longer have to stare at a fixed location.

When dynabooks become available, people will actually want to read material on screen rather than on paper. This will lead to yet another huge jump in the demand for online material.

We will need to have new forms of interaction when mobile computing becomes commonplace. Typing into a keyboard is not very convenient when you are driving down the freeway. (Unless it is Highway 101, that cuts through Silicon Valley---most of the time the traffic on 101 is sitting still anyway.)

It is likely that voice interaction will be the likely solution. Computers will speak to us and we will speak to them. Even now it is very simple to design a systems that will recognize a limited vocabulary from a single user. Recognition of multiple voices is not too far off. Such systems will allow for much more interaction, especially when mobile.

5.2.3 Filters and recommenders.

Given the forthcoming glut of information I expect to see widespread use of filtering systems and recommender systems. Filtering systems can be used to filter email, information, etc. based on the user’s preferences; recommenders help the user find interesting material. Such systems are not in widespread use now, but I expect this to change significantly.

One interesting economic aspect of social recommending systems such as GroupLens is that they have natural economies of scale. The larger the system, the more likely it is that you will be able to find someone ‘‘like you.’’ This means that it is likely that there will be only one or two large recommender systems for each topic (games, hobbies, reading, etc.) However, it seems unlikely that providers can extract much monopoly rent from such systems since entry is relatively easy. One likely business model is that they will be advertiser supported: a query will bring up user recommendations, along with commercial suggestions for related goods. For example, a query about good places to ski, may be displayed along with ads from ski manufacturers. Several Web search engines are currently experimenting with a similar model.

5.3 Policy

In terms of the network infrastructure, I think that it is likely that we will see some movement towards re-regulation. This is for two reasons. First, the tendency towards end-to-end service will require some form of antitrust oversight. Second, the network will be so important to our business and personal lives that it will *have* to be regulated. Just as gasoline prices are currently politicized, the Internet will be politicized. One can welcome this or bemoan it, but it is, I maintain, the likely outcome.

Another issue that will be highly relevant in they year 2000 is cybercrime. They once asked Willy Sutton why he robbed banks. His reply was ‘‘Because that’s where the money is.'' If there is a lot of money online, you can be sure that people will try to get to it. The commercial protocols themselves are likely to be quite secure due to encryption technology. There are two weak links: human error and the physical infrastructure.

User error can be dealt with by adding redundancy to the system. Smart cards and biometric information such as fingerprints or retinal prints are two obvious examples. Errors such as trusting the wrong person will always be with us, though systematizing the ''web of trust'' and adopting ''paranoid procedures'' may help control this sort of problem.

The physical infrastructure is more problematic. As more and more commercial activity moves onto the Internet, the entire economy becomes more and more vulnerable to terrorist activity. When you see how much damage a misguided backhoe can do, one has to worry about *intentional* sabotage.

By the year 2000, encryption will be widely used. In fact their will likely be several layers of encryption. Packets can be encrypted at the TCP level (within each user's machine) as well as within each application. Such widespread encryption will make commercial transactions reasonably secure, as well as alleviate some sorts of privacy issues---e.g., snooping.

However, other sorts of privacy issues will still be debated. The line between public information and private information can be drawn in many places, and exactly where that line can and will be constructed will require significant social debate.

6 Summary

The Internet is a mature technology, but it is an immature economy. Many of the fundamental legal, economic, organizational issues remain to be solved. The biggest danger, in my opinion, lies in making critical choices prematurely. There are many issues where the ''right'' answers are very obscure, and it will be easy for a poorly-informed or poorly-advised Congress to make the ''wrong'' choice. The critical issue should be to maintain the flexibility and openness that has made the Internet such a successful and vibrant phenomenon.