How the Internet Became Commercial

Innovation, Privatization, and the Birth of a New Network

Shane Greenstein

Princeton University Press
Princeton and Oxford

8

Overcoming Two Conundrums

A good predictor of not finding an ISP is the presence of hybrid corn seed.

—Zvi Griliches, Paul M. Warbug Professor of Economics, Harvard University[1]

What does the spread of the Internet have in common with the spread of hybrid corn seed? On the surface, there is not much similarity. Each episode is drawn from rather different eras and distinct parts of the economy. Below the surface of the merely superficial, however, the deep insights of Zvi Griliches illuminate the hidden forces shaping both.

Griliches was an economic giant of the twentieth century, and his biography made his accomplishments even more impressive. Born in Lithuania, the Nazi invasion shattered Griliches's youth, and in 1941 his family was moved into Kovno, the Jewish Ghetto.[2] In 1944 he and his family were interned near Dachau. A homeless orphan after liberation, Griliches immigrated to British Palestine, which then became the young state of Israel. Lacking any formal education as a teenager, and largely self-taught, Griliches enrolled for a short stint in Israeli schools for young Holocaust

[1]Zvi Griliches, private conversation, December 1996. He stated this after being shown the maps of the commercial dial-up ISP networks in the United States. This map is figure 8.1 in this chapter. An account of this remark and its context can be found in Greenstein (2004).

[2]For summaries of his contribution to economics and personal reflections on Griliches's contribution, see, e.g., Trajtenberg and Berndt (2001), 97. "There is something extremely rare, indeed unique about Zvi's overall package of human qualities—that combination of paternal warmth, intuitive understanding, ultimate wisdom, and cautious but firmly entrenched optimism."

survivors. Soon he gained permission to attend university in the United States. He paid for his undergraduate education with an unlikely inheritance—selling land his parents had bought many years earlier in British Palestine to help the Zionist movement.

Aspiring to become an agricultural economist, Griliches began his undergraduate studies at the University of California, Berkeley. His brilliance quickly became recognized, and a few years later he began his doctoral studies at the University of Chicago. There he completed his widely influential dissertation, "Hybrid Corn: An Exploration in the Economics of Technological Change," in which he demonstrates how to statistically measure the determinants of the adoption and diffusion of the hybrid corn seed, a tour de force for the time.[3] Upon finishing, the University of Chicago's Department of Economics hired him as a faculty member. In 1969, Griliches moved to Harvard, where he remained for the rest of his life.

Griliches's dissertation was more than merely an impressive statistical exercise. Griliches chose hybrid corn for study because it was one example of the productive advances transforming agriculture in the local and national agricultural economy. In just a few decades, the entire agricultural economy had become dramatically more productive, squeezing a great deal more out of a great deal less. Prosperity had noticeably increased, giving sons and daughters much higher incomes than their parents. What had happened in hybrid corn, in other words, had happened across a large range of agriculture.

As unlikely as it seems on the surface, the economic forces that shaped hybrid corn seed were similar to those shaping the diffusion of the Internet. In both examples widespread economic growth could not arise until many market participants figured out how to make each innovation inexpensive and easy to deploy in a wide set of circumstances. Each innovation had to overcome the *adaptation conundrum*. To become useful on a large scale, the innovation needed to become valuable in nearly every circumstance—namely, in every location, with every supplier, and with every set of users.

Consider the adaptation conundrum overcome by hybrid corn. Existing seed became less expensive and more productive, allowing farmers to

[3]See, e.g., Trajtenberg and Berndt (2001), 93. "Griliches's work on the diffusion of hybrid corn (the first article, stemming from his PhD dissertation, was published in 1957), catapulted him to instant fame, and has since turned into a true 'classic.'"

redirect resources to other activities. Better seed resulted in more output (better corn per stalk) from all producers, and more output per input (land, labor, fertilizer) from any individual producer. That freed up spare time and resources for other activities. Fertilizers and pesticides eventually changed to accommodate a new scale and mix of crops. Crop failures became less likely, which altered the financial risks. Bigger equipment could be purchased for the more productive land. New uses were found for abundant and inexpensive corn. Seed vendors altered their inventories, carrying different supplies. Harvesting equipment designs evolved; therefore, the manufacturing stamps, tools, and dye changed too. Finally, new books about best practice were written, and educational practices were adjusted. Agricultural extension stations taught those methods, and they passed between farmers by word of mouth.

Said succinctly, numerous *work processes*, *organizational routines*, and *market relationships* were restructured to take advantage of new potential and new opportunities. No general economic law guaranteed whether and when such adaptation took place, or, for that matter, whether it succeeded at all. Widespread success required inventiveness from many market participants. The innovation had to be adapted to a multitude of circumstances faced by potential suppliers and buyers, and the hindrances and bottlenecks found in every nook and cranny of every region had to be overcome. Almost by definition, every inventive act took place at a unique point in a technology's development, and often the market circumstances were novel also.

As this chapter and the next will explain, the Internet would not spread and lead to economic growth until it too overcame the adaptation conundrum. Such growth involved inventive acts from many market participants—vendors and users, and their business partners and consultants.

The Internet and corn differ in one important respect. The Internet had to overcome an additional conundrum, the *circular conundrum*. In technically advanced communications networks, the circular conundrum arises because of the large number of economic actors providing technically interdependent services. Such interdependence makes them complements in demand, meaning that the maximum value of using any service cannot be achieved unless all the components and services are used together. The mere presence of complements in demand does not guarantee the coordinated supplier of all the components, however. If anything, collective

fence sitting is the most common response to the circular conundrum—a response that manifests itself as cautious investment behavior by a group of erstwhile business partners and competitors.

Caution was common in 1994 and a real hindrance to the growth of the commercial Internet. A retailer might aspire to sell goods electronically, but would hesitate to spend for an inventive new website unless there were enough browser users it could reach. Or, a bulletin board firm might hesitate to provide Internet access in a novel location unless enough local businesses and users had computers equipped to handle the web. A computer consultant might delay in investing the time to learn HTML and adapt it to her client's needs unless she was sure enough clients had customers who would make use of electronic commerce. Or, a system administrator might pause before spending to install novel software to support browsers and gateways for electronic mail unless it were obvious that employees were going to use the new services. Each participant remained cautious, waiting for someone else to go first.

Starting in 1995 and accelerating thereafter, the US economy overcame the circular conundrum at the same time it overcame the adaptation conundrum. It is not possible to distinguish when one was overcome from when the other was overcome.

In the mid-1990s, there were several areas to observe triumphs over the two conundrums, but in this chapter, I focus on one key area, specifically, ISPs.[4] ISPs are a good illustration because they were an entirely new business, created for the commercial Internet. The most established firms in the ISP market, such as MCI, Sprint, and BBN, had fallen into the business after their involvement with the NSF. But the many more entrepreneurial firms that came from the edges of computing and communications, such as PSINet and UUNET, also played an essential role.

The behavior of ISPs in late 1996 motivated Griliches to make the quip quoted at the top of the chapter about the relationship between hybrid corn and ISPs. His observation occurred at a moment when the two conundrums had *not entirely* been overcome. There was virtually no Internet access in the places where corn grew, or, for that matter, in US counties where the cows and cotton outnumbered the local inhabitants. At that point, the commercial

[4]The next chapter focuses on another key area—providers of Internet and web services to business.

Internet aimed almost exclusively at rich technological enthusiasts in major urban areas. In two short years, however, the situation changed. It would become clear that technology enthusiasts would no longer be the sole users and, correspondingly, suppliers could reach rural areas. Commercial markets could supply the Internet nearly everywhere and in forms that appealed to a large number of nontechnical households and businesses.

Building and Adapting the Internet[5]

In 1994, it appeared as if there was only one known solution to the circular conundrum. It involved two risky steps, neither of which could occur quickly.

- First step: An important economic actor had to take on considerable risk by investing ahead of (a) others, (b) all other suppliers, and, crucially, (c) realized demand.
- Second step: An early mover had to motivate other vendors to follow with more investment.

Those two steps needed to occur before most mainstream suppliers would act. Only after others had taken the risks would the entire set of suppliers learn whether their biggest ambitions would be realized—namely, whether there would be enough revenue to pay back the collective expense from entering the market.

Before the Internet grew, however, most established firms believed only one type of economic actor could take that first step—a very large established firm that could financially handle a loss if a big risk failed, and that possessed the ability to bring others along with it. And only a few large established firms with recognizable brands could have played that role—namely, IBM, Intel, Microsoft, or AT&T, for example. Crucially, and each for different reasons, none of these firms had publicly backed the Internet as a technology for mass-market use prior to 1995, and, accordingly, many other firms hesitated to invest.

[5]The themes for this chapter have antecedents in research about the diffusion of general purpose technologies (e.g., Bresnahan and Trajtenberg 1995) and *coinvention* (e.g., Bresnahan and Greenstein 1997, 1999). Coinvention is the invention complementary to the general purpose technology, which makes it useful in a variety of settings. Coinvention can take place at user installations, vendor organizations, and the establishments of other participants in the economy around a general purpose technology. The themes in this chapter borrow extensively from writing about coinvention at ISPs (Greenstein 2000a, 2000b, 2001, 2005, 2008b, 2010b).

The conventional view turned out to be wrong because it remained blind to one other possibility—that a group of inventors and suppliers could collectively build a technology good enough to deliver value for many users. That too involved a slow two-step process, building first for technical enthusiasts and then later for general users.

As prior chapters have discussed, the potential to build for general users had eluded others. Events in the Internet illustrate why: Netscape sold an easy-to-use browser, building on tools developed by Tim Berners-Lee and operating on a network the NSF had just privatized. The multifaceted combination of the NSF's investments and privatization, many ISP's investments and delivery, and browsers built in the web led to a feasible network. And where the young Netscape led, many firms had to follow with real investment, anticipating enough demand for their services to justify the expenses they incurred.

There was more, as well. Between 1994 and 1995, a number of other young firms also began exploring new businesses built with web software, including Yahoo!, eBay, Amazon, Vermeer, and others. In addition, established firms began experimenting with web pages, including Sony, *Encyclopædia Britannica*, and others. With the release of the Netscape browser, and especially after the success of the Netscape IPO, numerous software entrepreneurs began making plans for new start-up businesses, and these plans did not depend on whether a large firm took action or not. In other words, the actions of AT&T, IBM, or Microsoft did not play a role.

As earlier chapters described, Netscape's and Microsoft's growing rivalry would eventually contribute to moving this entire market forward for a while. From the outset, both priced their browsers at an inexpensive point. At first Netscape allowed home users to download their browser for free and charged businesses to use a license. In comparison to typical enterprise software in most establishments, these charges were rather minimal.

The biggest expenses for ISPs were the human costs of installing the software and changing operating procedures, and those expenses were also much lower than typical. The Netscape browser was easy to install, and so were some of the basic tools for managing users. Netscape crowed about its compatibility with other TCP/IP and web software, encouraging users to install suites of such applications. Microsoft's Windows 95 also contributed to making it easier to go online. It included built-in TCP/IP

compatibility, and, importantly, its networking software, Windows NT, eased the transition for many businesses.

After December 1995, prices sank even lower. Microsoft began pricing its browser at zero, bundling it with every operating system it sold. That made the browser nominally inexpensive, by definition, albeit, not any less burdensome to install. More important, Microsoft (1) began announcing plans to integrate its browser with its server software and other application software and (2) aggressively began to make it easier for businesses to use its browser. These events in browsers would have consequences for other parts of the young value chain for Internet services, because this crucial gateway to the Internet became seamless to users.

Why Internet Service Providers Mattered

To a user, there were two costs. The first were the hardware costs. The second involved the human costs, the hassles affiliated with learning how to use the software and adapt it to the unique needs of the user or business establishment. By 1996 installation costs became the principal hardware cost of adopting the Internet and World Wide Web. Installation meant establishing a connection with an ISP and adapting the connection to an existing PC or local area network (LAN).

The actions of ISPs played a significant role at this point. These were a new type of firm, and, after the browser firms, the most visible suppliers of the Internet and World Wide Web. Many aspects of the commercial Internet's value arose from the decisions of ISPs. These decisions determined how much many users valued the Internet, because many ISPs determined the appearance, packaging, and pricing of the initial Internet experience for many users.

By 1996, ISPs offered service in every major US city, and many large firms had begun building national networks. How did that happen? No dramatic technological invention lay behind these changes. Rather, a new and stable value chain emerged. It came about due to the interaction of three related inventions aimed squarely at the adaptation conundrum:

- Invention of a new pricing norm;
- Invention of a viable model for distributing Internet access over geographic space;
- Invention of new points of differentiation between ISPs.

Before discussing each of these in detail, consider why the overall market was so nurturing.

Markets for business access operated on straightforward principles. Sprint, PSINet, MCI, and plenty of other firms had figured out the value chain. Business establishments asked for either fixed-line access or dial-up access. Fixed-line access involved delivering Internet connections of various capacities in various locations while also interacting with an Internet gateway. Norms in that business followed administrative precedents set in the research network. Establishments controlled their data on one side of their own gateway, usually because the gateway lay at the end of a LAN. The ISPs controlled everything after the data was handed off from the gateway.

Fixed-line access involved figuring out how to get physical presence near customers. These challenges could be regulatory, as in gaining rights of way. The challenges also could be mundane, as in negotiating with various builders in order to snake underground lines throughout a city, which was expensive and not typically desirable. In either case, the challenges were not mysterious. ISPs knew what challenges needed attention in order to be in the business.

Dial-up access faced different challenges. The biggest puzzles were associated with serving unsophisticated users. Many, if not most, of the technological enthusiasts or wild ducks had adopted by 1995. After 1995 the users fit a different profile. They wanted a low-risk valuable service, one that was reliable, routine, and inexpensive. But attempts to reach unsophisticated users in the first half of the 1990s had made only marginal progress. Whereas AOL, CompuServe, Prodigy, and Genie had no more than several million users between them, PSINet and UUNET largely focused on business users and avoided home markets. NetCom and Spry did manage to build a customer base in the hundreds of thousands, but they were notable for their rarity as no other ISPs were getting the same scale of success. If there was a large untapped market, nobody was quite sure how to tap it.[6]

[6]The biggest problem was not with the technology per se. A number of other routine technical and business processes had begun to characterize the operations of ISPs. Technical practices improved. For a broader discussion of these and other technical issues, see Haigh (2008).

The invention and spread of the World Wide Web—namely, three spreading software tools, HTTP, HTML, and the URL, opened a new range of possibilities for nontechnical users. In addition, invention of the web enabled the growth of such a new potential market beyond e-mail, one that involved pictures and graphics. That potential was widely recognized among virtually every ISP by mid-1995. Nonetheless, it did not bring millions of users to the doors of ISPs. Somebody still had to show users that the technology did something useful and was worth paying for. In the meantime, ISPs had to price the service to cover their costs and appeal to users.

Pricing Internet Access

Pricing had vexed firms for years. In the first half of the 1990s, most ISPs tried the same pricing norms governing bulletin boards. For bulletin boards the pricing structure of the majority of services involved a subscription charge (on a monthly or yearly basis) *and* an hourly fee for usage. Some prices for specific services were unlimited, while others had usage charges, and different firms experimented with different combinations.[7] For many applications, users could get online for "bursts" of time, which would reduce the total size of usage fees. Although it was cumbersome to log on and off, technical users did not feel hindered. E-mail, file transfer, and posting to listservs could accommodate these practices. But nontechnical users resisted these logging procedures as burdensome, and most did not think the benefits worth the hassles.

The emergence of faster and cheaper modems and large-scale modem banks with lower per-port costs opened the possibility for a different pricing norm, one that did not minimize the time users spent on the telephone communicating with a server. Modems allowed a user to move data over a digital phone line, at this time typically at a maximum speed of 14.4 and 28.8 kilobits per second. Modem banks allowed an ISP to handle multiple phone calls. The emergence of low-cost routines for accessing a massive number of phone lines also contributed to a new norm, because it enabled many ISPs to set up modem banks at a scale only rarely seen during the bulletin board era.

[7]Banks (2008) provides a long history of many of these experiments prior to the rise of the Internet.

That did not mean a successful solution was initially obvious to many vendors in 1995. Judging by the advertisements they used in *Boardwatch Magazine*, solutions varied widely. Some ISPs tried one thing, while others tried another, symptomatic that no consensus emerged right away.[8] Although there was a long-term trend toward more use of unlimited prices, the national ISPs also tried a variety of usage charges.

Two broad viewpoints emerged. One viewpoint paid close attention to user complaints about monitoring. Users remarked that surfing the World Wide Web was hypnotic. Users found it challenging to monitor time online. In addition, monitoring time online was nearly impossible with multiple users within one household. The ISPs with sympathy for these user complaints argued that unlimited usage for a fixed monthly price would be a feasible solution. A premium would cover the extra costs of the unlimited usage, though it was an open question how large that premium needed to be. These plans were commonly referred to as *flat rate* or *unlimited* plans.

An opposite viewpoint contrasted with the first one. Specifically, user complaints were transitory, and new users had to be "trained" to monitor their own time online. Supporters of this view pointed to cellular telephones and bulletin boards as examples. Cellular telephony was also beginning to grow, and pricing per minute had not deterred users. One up-and-coming bulletin board firm, AOL, had seemed to grow with such usage pricing as well. The ISPs that held to this view expressed confidence that usage prices would prevail, and they called for exploration of new combinations that might better suit surfing behavior by nontechnical users.

While the events that motivated these two viewpoints were new, the outlines behind them were not. The disagreement reflected a long-standing argument about the appropriate method for pricing telephone services. Each side could burnish its arguments with considerable experience from the past, and the argument could have continued indefinitely.

Although both sides recognized that the Internet might break with precedent, a competitive event settled the argument among ISPs, and a positive response from users reinforced the result. The key action came from AT&T, one of the established firms that could move a market. AT&T

[8]For example, the advertisements in *Boardwatch Magazine* showed such a variety. See Stranger and Greenstein (2006, 2007).

did move the market, in fact, but just not as AT&T's executives had expected.

Pricing at a New ISP

AT&T's Worldnet service was first aimed at business in late 1995 and then explicitly marketed at households in early 1996. Held in low regard by many technically oriented participants in the Internet, AT&T was not expected to succeed. When their success went well beyond expectations, it caught a lot of attention and framed a simple question: How was this possible? Although flat rate pricing was not solely responsible for AT&T's initial success, it was one of the unusual features of AT&T's service. Their success led other ISPs to consider imitating AT&T's *approach* to flat rate pricing.

Viewed as a potential competitor to virtually every ISP, the competitive response was swift and almost complete. Virtually all ISPs switched to unlimited usage except AOL and Microsoft Network (MSN). Why did AT&T's action induce such a swift response from all but these two? AT&T's strategy was aimed squarely at the nontechnical user at home, the same user community that all these firms were trying to court. AOL's holdout requires explanation, which in turn requires understanding AT&T's actions.

AT&T took a calculated business risk by choosing a design for its service that would not—indeed, could not—make a profit if it did not get widely adopted in many cities throughout the United States. The cost structure for AT&T placed challenging constraints on the network's design. AT&T aspired to make the service both national and inexpensive but could not meet those aspirations unless it achieved large scale in its operations. If it did achieve such scale, then its costs per user would be low, allowing AT&T to charge close to the lowest prices in the industry at that time. If it did not achieve such scale, then the costs would be too high, and the service would be unprofitable.

How would AT&T get so many users? AT&T's executives saw an opportunity to be among the first to offer home service that came from a branded and established national company, which appealed to certain users who did not trust young or unfamiliar firms. The executives compared AT&T with another potential entrant with similar credentials, IBM. Although IBM supported a worldwide service for business customers, the company had made it clear in 1994 that it did not aspire to provide service to the mass-market home user.

One part of AT&T's strategy was built on the path forged by MCI, Sprint, UNNET, and PSINet, which already had begun to demonstrate how to offer service as a national firm. A national firm needed a geographically dispersed set of investments because virtually all users preferred to access the Internet using local telephone calls. The US telephone universal service policy had long subsidized local landline household calls by charging greater prices for business and long-distance calls than for local calls, which were priced at a flat rate per month in almost every state, with the definition of *local* left to a state regulator. Typically *local* was defined over a radius of ten to fifteen miles, sometimes less in dense urban areas, such as Manhattan. Thus using local rates reduced household user expenses, thereby making the service more attractive.

If AT&T sought to provide service for an entire city, its management had to assemble a list of local phone numbers to permit all potential callers to have one or more options to make a local phone call. Providing national service meant assembling local numbers across many cities and states.[9]

As it turned out, with between four or five hundred phone numbers, AT&T, as well as any other ISP, could cover every major urban area of the United States—namely, any city with a population over fifty thousand inhabitants. Approximately 85 percent of the US population lived in such cities. In brief, with such a network AT&T could claim it had a national service.

A national network, however, involved more than phone numbers, which is where ISP service became expensive. Each phone number needed to be outfitted as a point-of-presence, or POP, a term borrowed from the NSFNET. The modem banks that handled the calls at POPs were one of the biggest expenses of an ISP.

Attracting nontechnical users was a strategic plan AT&T used to spread the expense of modem banks over a larger base of users. Their strategy bet that nontechnical users would *not* stay online nearly as long as experienced or technically oriented users—in effect, anywhere from twenty to thirty hours of time a month at most, and often a lot less. Because such light users did not all dial-in on the same day or at the same time of day,

[9]The US telephone network consisted of approximately 19,000 local telephone switches and each of those switches was affiliated with a single area code and multiple prefixes. For example, in the phone number 312-555-1212, the area code is 312 and the prefix 555. To offer dial-up Internet service in a city, AT&T had to compile a list of the local area codes and prefixes that covered the area, and then implement its service. Such information was readily available.

the equipment investment did not need to handle all calls at once. Rather, with light users, an ISP could serve a local area with modem bank capacity at anywhere from one-third to one-quarter the size of the total local service base.

More to the point, with such ratios, AT&T could offer the service at a low price. Their executives merely had to find the lowest price most ISPs offered at the time. The marketing department determined $20 per month, which generated revenue of $240 dollars per year per customer, which, in turn, was enough to cover the investment and operating cost of a POP with far less than one thousand customers.

Those cost and revenue scenarios made the service viable in many medium-sized cities. For example, a city of one hundred thousand inhabitants typically contained forty thousand households. If one-quarter of such households had personal computers at home (a plausible number for 1996), then there were ten thousand potential customers. To AT&T's marketing executives, gaining 10 percent market share in such a city seemed obtainable. And if that was obtainable in a town of one hundred thousand, then similar results seemed plausible in bigger cities, like Chicago, Atlanta, and Denver. It also meant that the financial risk of offering service in smaller cities was not too great, so AT&T targeted cities of fifty thousand inhabitants, even in isolated locations.

AT&T unveiled its home service and in a few months it signed up one million customers. This experience was stunning. It was the type of counterexample that attracted the attention of every ISP in the country. Alarmists among longtime Internet watchers wondered if AT&T, the firm many wild ducks resented and despised, would dominate home markets for Internet access. Virtually all small and medium ISPs quickly matched the $20 price and offered unlimited service as an option. Many tried to offer an additional option, usage pricing with some types of discounts, but over time many would drop the practice due to lack of customer interest.

The Response at AOL

Did AT&T's actions and its competitor's reactions settle the pricing norms for good? Not in the least. In early 1996 only two major ISPs, AOL and MSN, refused to change their pricing practices quickly. AOL was particularly interesting because it had one of the largest user bases in the country, so its hesitance received considerable attention.

Inheriting usage pricing from its long legacy as a bulletin board firm, AOL aspired to appeal to the nontechnical users with superior marketing, appealing page design, community building in moderated forums, and tailored content offerings, not just branding or technological prowess. Its existing pricing had worked well, covering the operating expenses. And a pricing change had not been part of its plans.

As ISPs converted to flat rate pricing, AOL's management found itself between a rock and a hard place. Multiple considerations hemmed in its options. If it did not change its pricing, AOL's management feared it would increasingly lose potential customers to other ISPs, especially national ones. Like other ISPs, AOL perceived the potential for growth by attracting new users over the next years, so choosing the wrong pricing could be a very costly mistake during the earliest moments when many new users came online. Yet, moving to flat rate pricing presented many immediate challenges. In 1996, AOL owned and managed its own facilities for providing access. A new pricing policy would undo years of fine tuning of capacity in multiple locales, which was designed to match historical (and tested) ratios between the number of local subscribers and the capacity needed to handle them.

This last observation was especially salient. AOL's users did not overuse modem capacity, because a large number avoided charges by logging in for short periods of time, thereby freeing up modem capacity. AOL's executives guessed (sensibly) that flat rate pricing would lead existing users to stay online longer than with usage prices. They also anticipated their capacity of phone numbers and modem banks in many locales would be inadequate, as would server capacity in some locations.

The competitive pressures were too much, however, and the potential loss of customers to MSN forced a change,[10] and AOL announced that it would match flat rate pricing, taking effect December 1, 1996. Capacity is-

[10]At the time MSN also did not have flat rate pricing. MSN's announcement that it would switch became the final catalyst for AOL to make the change. Steve Case, AOL CEO, reminisced many years after the change that AOL had studied the potential switch for quite some time, but had not acted on it because management could anticipate a difficult transition.

> It came to a head over a weekend as Microsoft announced they were offering MSN on a flat rate basis, and it was clear they were planning to steal a lot of market share from AOL. So I decided within hours of their announcement that we had to match them, and the company worked throughout a weekend so we could make an announcement.

sues did emerge in many locales, manifesting as busy signals, just as AOL's executives had feared.[11] These well-publicized troubles continued for many months, earning AOL a poor reputation among some users. The bad publicity induced further entry by other ISPs looking to acquire customers fleeing the busy phone lines. Ultimately, AOL survived the bad publicity through a series of new investments in facilities, new arrangement for content, and more intense marketing—that is, carpet bombing of trial CDs. That was not a foregone conclusion throughout most of early 1997.

This experience led AOL to alter many aspects of its business. It positioned itself as "the Internet and a whole lot more."[12] In the following year AOL sold off its facilities for access and contracted with others to manage them, while also acquiring the customers of CompuServe. While it became a software company primarily, its software became the most common gateway to the web in the United States in the late 1990s.

Tensions with the Pricing Model

Throughout 1996, 1997, and 1998, ISPs experimented with hourly limits and high marginal pricing above the limit. These experiments largely failed, and those failures demonstrated how binding the unlimited contract became as a norm for pricing. The experiments responded to a stark problem: Many technical users, and even some nontechnical users, stayed online for more than twenty to thirty hours a month, ruining the key economics that allowed ISPs to share equipment across many users.

High usage could happen for a variety of reasons. For example, some technical users simply enjoyed being online for large lengths of time, surfing the growing Internet and web. Some users began to operate businesses from their homes, remaining online throughout the entire workday. Some users simply forgot to log off, leaving their computers running and tying up the telephone line supporting the connection to the PC. And some users grew more experienced and found a vast array of activities more attractive over time. The list of reasons went on and on.

See http://www.quora.com/AOL/How-did-AOL-make-the-decision-to-go-to-an-all-you-can-eat-pricing-strategy/, accessed November 2012.

[11] Less than three weeks later a company spokesman said seven million AOL customers in the United States were conducting 30 percent more daily online sessions, lasting an average of 20 percent longer. San Jose Mercury News (1996).

[12] See http://www.quora.com/AOL/How-did-AOL-make-the-decision-to-go-to-an-all-you-can-eat-pricing-strategy/, accessed November 2012.

Many ISPs tried to offer limited contracts in addition to their unlimited contracts, discounting the price (from the unlimited benchmark) in exchange for the limit. Most such limits were not particularly binding—involving monthly limits ranging from sixty to one hundred hours—unless the user remained online for hours at a time most days of the month. Some ISPs tried offering steep discounts for steep limits, such as $10 discounts for thirty hours a month. Yet few buyers took them, persisting with the slightly more expensive unlimited contracts.

A number of ISPs also included fine print that allowed them to automatically log off users after a large number of consecutive hours, or after their accounts remained inactive for a fixed period, such as thirty minutes. Because some customers were unaware of or did not read the fine print, they perceived such types of automatic logging off as poor service or unreliable equipment from the supplier. Consequently, many small ISPs hesitated to employ this approach.

A few ISPs deliberately adopted contracting policies that tried to move the costliest users by refusing to offer unlimited contracts and thereby forcing all users into hourly limits per month. Such limits imposed costs on only the most intensive users, which was a strategy to induce their most expensive customers to migrate to another vendor. If an ISP already had a strong customer base, the loss of a few intensive and costly users was perceived as a gain.

As in urban or high-density environments, limits also arose in settings without many suppliers, such as low-density environments. Some ISPs adopted limits because there was no substitute and this was the simplest way for them to limit the expensive user's activity. Because such actions were more common among ISPs in low-density environments, it suggested the practice could survive more easily in the face of less competition.

Despite all these experiments, the unlimited contract dominated transactions. One survey of pricing contracts in May 1996 found that nearly 75 percent of the ISPs offering 28K service (the maximum dial-up speed at the time) offered a limited plan in addition to their unlimited plan. That dropped to nearly 50 percent by August. By March 1997 it was 33 percent, 25 percent by January 1998, and less than 15 percent by January 1999.[13]

[13]See Stranger and Greenstein (2006, 2007).

Overall, why did unlimited prices become the norm? Once the market moved to that general practice, ISPs could not find many takers for the alternatives. Competitive forces kept matters oriented toward the user's preferences for unlimited capacity. Demand and competition forced suppliers to adapt, and ISPs had to figure out how to manage their businesses to support that preference.

Geographic Spreading

Not all ISPs were alike. Different firms came to different conclusions about how much territory to cover. Broadly speaking, ISPs could be divided into three groups, backbone providers, national access providers, and local access providers. Before 1996 it was unclear how these would differ. After the emergence of a pricing norm, all used the same basic economies of POPs. Nonetheless, these three types of firms behaved differently.

The first group—the backbone providers—was private national firms (that is, MCI, Sprint, UUNET, BBN), the largest carriers of data, and the followers of the NSF's original plans. In 1995 and 1996, all of them carried traffic from NSF's NAPs and the CIX. Any regional ISP could exchange traffic with them. This aspect of backbone behavior would begin to change in 1997, but not at this crucial moment, when the ISPs were trying to overcome the adaptation and circular conundrums. At this point, the backbone of the US Internet resembled a mesh, with every large firm both interconnecting with every other and exchanging traffic with smaller firms, just as NSF had conceived of it by fostering many NAPs.

The mesh carried from the backbone to the retail level. Many of these same national firms also provided retail ISP services to consumers or to other ISPs that rented rights to resell use of their modem banks. Not all took the same approach, but in a mesh it did not mater. Some of these firms owned their own fiber (for example, MCI) and some of them ran their backbones on fiber rented from others (for example, UUNET).

The next two groups, the national and local providers, were ISPs that ranged in size and scale from wholesale regional firms down to the local ISP handling a small number of dial-in customers. Many of the largest firms developed familiar brand names, such as EarthLink, Sprint, AT&T, IBM Global Network, AOL, and MindSpring. Other large firms included entrants or Internet insiders from the NSF, such as PSINet, Netcom, ANS, and GTE (which in 1997 acquired assets from BBN [Bolt, Beranek and

Newman] Planet) and others. Still others offered ISP services to consumers (for example, AOL and MSN), and might not own any facilities. They rented it from others (such as ANS), though users did not necessarily know this.

Market share at the retail level was skewed. A couple dozen of the largest firms accounted for 75 percent of market share nationally and a couple hundred for 90 percent of market share. In other words, the majority of these ISPs were small dial-ups covering a small regional area, but the majority of users employed national providers.[14]

A large number of small ISPs began to flood the market in 1996. The rapid growth of ISPs seemed almost magical to users with no BBS or Internet experience, as ISPs seemed to pop up instantly out of nowhere, and in every major city. For example, in May 1996 *Boardwatch* listed prices for just over two thousand ISPs. In August it listed prices for 2,934 ISPs. The proliferation of ISPs continued. In February 1997 *Boardwatch* listed prices for 3,535, in January 1998 for 4,167, and in January 1999 for 4,511. In each case, though, the magazine listed many more ISPs for whom they did not have price information. The highest reported number in *Boardwatch* was just over seven thousand in March 2000.[15] Half of these ISPs supported only one phone number, and the vast majority of the rest supported anywhere from two to six hundred numbers. The vast majority of ISPs offered unlimited service.

Coincident with this growth, Internet access became widespread. Notably, providing Internet access was feasible to almost anyone with technical skill, especially former BBS operators, many of whom wanted to open a POP in uncovered territory. Many entrepreneurs perceived that AT&T and AOL did not appeal to every user, and a local ISP could tailor its services to local customer needs. The economies of POPs allowed ISPs to survive with comparatively small customer bases, serving limited geographic customer bases. Many small ISPs ran their own network POPs and provided limited geographic coverage.

Also, many small ISPs could cover their costs on a small scale. Moreover, few barriers existed for an ISP. The standard server software for

[14]The National Telecommunications and Information Administration (NTIA) surveys do not begin to track broadband users until August 2000, when the survey finds that 4.4 percent of US households are broadband users (with 41.5 percent of households being Internet users).

[15]See, e.g., Downes and Greenstein (2002, 2007). Also see Stranger and Greenstein (2006).

supporting an ISP, Apache, was available without restriction, and the necessary technical know-how for getting started did not differ greatly from routine knowledge found at a firm performing computing services or bulletin board services prior to commercialization. Apache ran on Unix, which was increasingly becoming Linux, and many ISPs were operated by programmers who had the technical skills to manage this type of server.

The mesh of the backbone helped. Many small firms leased lines from larger ISPs, such as Sprint, or externally managed POPs through a third party, such as AT&T or MCI, in locations where they did not have coverage. This allowed them to offer national phone numbers to their local customers, which meant they were, in effect, competing with commercial online services with national coverage.

Figure 8.1 comes from fall of 1996. This was the map that prompted Zvi Griliches's remark about the Internet and hybrid corn, which was quoted at the outset of the chapter. There was never any issue about getting some service, even in the worse situation. A user could make a long-distance phone call and get service, but the effective cost was higher. In that sense most agricultural counties in the United States lacked local ISPs. As a rule of thumb, if corn, cotton, or cows dominated the county's economy in 1996 then the probability of finding an ISP in the local calling area was quite low.

Broadly speaking, therefore, the locations lacking local Internet access also were the poorest urban areas or the smallest and most remote rural populations. Some of these areas bore signs of permanent retardation of development, as in Appalachia or the Mississippi delta regions.[16] But more than poverty was at work in many locations. Rural areas often lacked a competitive supply of providers, and even when suppliers existed, they sometimes provided limited services or focused on specific segments, such as business users.[17]

The geographic spread of access changed, and it changed rapidly. A comparison of the fall of 1996 and 1998 captures these changes. In the fall of 1996, the entire commercial ISP network was estimated to be twelve thousand phone numbers for over three thousand ISPs. By the fall of 1998, the estimates were over sixty-five thousand phone numbers for just over six thousand

[16] See Strover, Oden, and Inagaki (2002). See also Strover (2001).
[17] See Nicholas (2000).

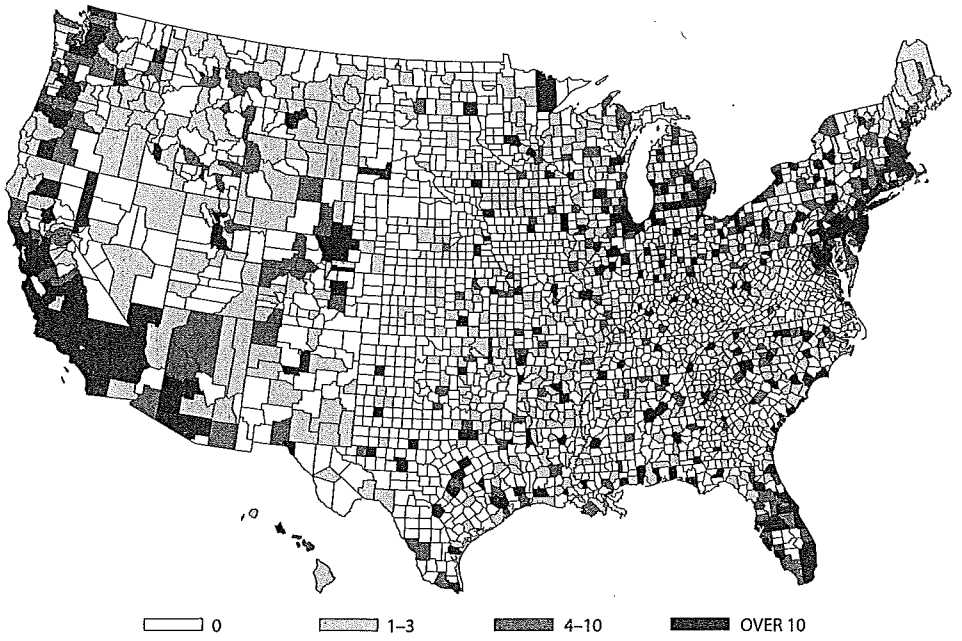

FIGURE 8.1. Distribution of ISPs, September 1996

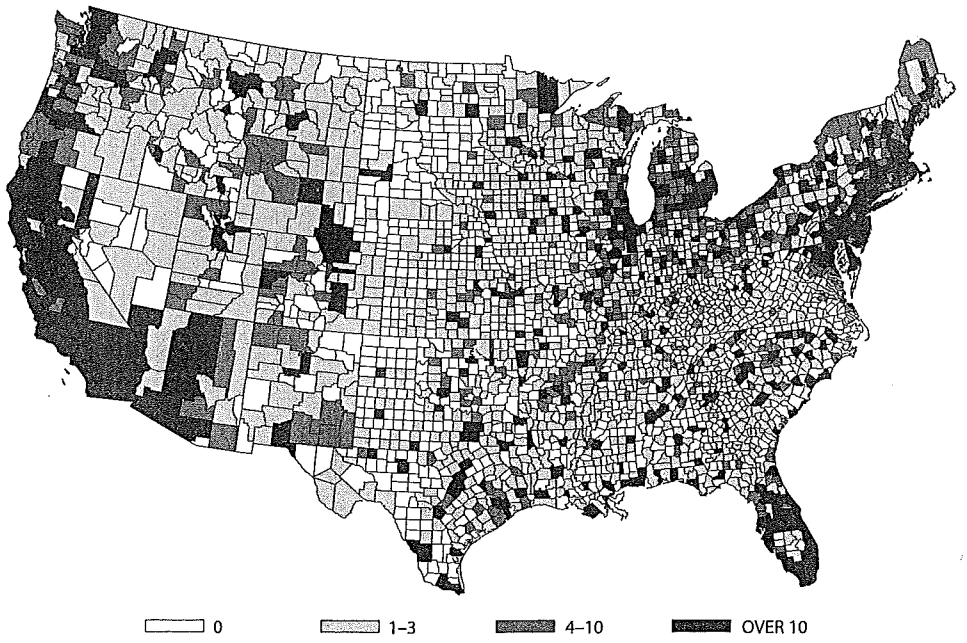

FIGURE 8.2. Distribution of ISPs, October 1998

ISPs.[18] In addition, one estimate showed that in 1998 more than 92 percent of the US population had access through a local phone call to seven or more ISPs, and less than 5 percent did not have any access.[19] Almost certainly, these estimates are conservative. The actual percentage of the population without access to a competitive dial-up market was much lower than 5 percent.

Figure 8.2 shows the situation in 1998. While many ISPs had locations in all the major population centers, there were also some providers located in sparsely populated rural areas.[20] The commercial Internet had acquired the capacity to support the delivery of the Internet to half the households in the country, to virtually every medium and large business establishment, and to a large fraction of small businesses.

In 1996, national firms found it in their interest to offer service with somewhere between four and six hundred phone numbers and only in populations of fifty thousand or above. In contrast, the vast majority of the coverage in rural areas came from local ISPs. For rural counties with a population under fifty thousand, the providers were overwhelmingly local or regional.

In the fall of 1998, the equivalent figures for national firms were thirty thousand or lower, which indicates that some national firms had moved into slightly smaller areas and less dense geographic locations.[21] Additionally, many rural telephone cooperatives were opening Internet service. This followed long-standing traditions in small and rural cities to use collective quasi-public organizations to provide utility services that other private firms did not find profitable.

There was a dichotomy between the growth patterns of entrepreneurial firms that became national and those that became regional. National firms grew by starting with major cities across the country and then progressively moving to cities of smaller populations. Firms with a regional focus grew into geographically contiguous areas, seemingly regardless of urban or rural features.

The stark differences between figures 8.1 and 8.2 showed that ISPs had begun to invent a subtle solution to the adaptation conundrum. The arrangement borrowed from the practices that existed at bulletin boards and

[18]See Downes and Greenstein (1998).

[19]See Downes and Greenstein (2002).

[20]See Downes and Greenstein (2002) for a description of the dial-up market, or Downes and Greenstein (2007) for an analysis for why some areas had more entry than others.

[21]See Downes and Greenstein (2002).

the NSFNET, and, crucially, could be mastered by any technically oriented entrepreneur with the time to learn them. Involving a subtle combination of inventions at the POP level, changes to pricing, and alterations in operations to support inexpensive pricing, this solution built on the mesh and incorporated a response to local conditions, which put the ISP in a position to mediate between groups of local customers, their telephone companies, and the national backbone firms.

Experiments in Differentiation

By 1997 many users began to take flat rate pricing for granted. To attract users, ISPs had to respond to competitive pressure to offer flat rate pricing and something else users found appealing. For local ISPs especially, the volumes of customers generated revenue to cover costs but did not yield much profitability. These ISPs looked for additional revenue by offering additional services to their customers of ISP service.

A crucial question concerned how the user employed his or her browser. From an ISP's perspective, the question shaped the design of the opening home page—or, as it was subsequently labeled, *portal*—that users would see when they first clicked on their browser.[22] An ISP had three potential choices. Should it (1) design its own portal potentially at great expense, (2) default to another's portal, such as Excite or Yahoo!, or (3) leave the decision to a user altogether?

Different ISPs made distinct choices and learned different lessons about the trade-offs between these choices. For example, some ISPs maintained minimal home pages, which many marketed as a virtuous attempt to give users freedom to choose for themselves among Yahoo!, Excite, Lycos, and a myriad of other young portals then springing up. Of these, a portion succeeded with—or in some views, in spite of—this choice.

The choice AOL made was unique, and while it may seem savvy in retrospect, many Internet enthusiasts then regarded it as risky. Seeking to attract the nontechnical user, and building on its experience prior the rise of the Internet, AOL invested in a *walled garden*. This approach protected content and tailored it to users, or, in the eyes of technically sophisticated vendors, "spoon fed" content to users. As part of this strategy, AOL would

[22]For a review of technologies behind the browser and portal, and an analysis of the factors that shaped their evolution, see Haigh (2008).

merge with young firms or make deals with certain firms to feature their service or copyrighted content, making it easily accessible to users.[23] That left AOL controlling a large fraction of the user experience while sacrificing sophisticated users to other ISPs.

Yet AOL's strategy also continued activity it already had performed in the era of bulletin boards. Management thought that AOL's prior investments in community building would continue to have value as its users transitioned to using the Internet more frequently, and AOL further supplemented those actions with other practices that some of its largest rivals, such as CompuServe, did not pursue, such as using log-in names and e-mail addresses using natural language labels instead of combinations of letters and numbers.

Not all of the ISPs' choices succeeded. For example, CompuServe, Prodigy, and Genie all failed at an approach similar to AOL's. The success of AOL's campaign was frequently attributed to its aggressive marketing—particularly the "carpet bombing" of its disks in news inserts and other unlikely places,[24] which facilitated a trial use of its service in households across the United States.[25] As an interesting contrast, MSN attempted a similar marketing strategy, and, with the help of its marketing advantages and budgetary tolerance for operating losses, did not fail. Nevertheless, MSN was no better than a distant second to AOL in market share throughout the 1990s.

[23]See Elfenbein and Lerner (2003).

[24]The aggressive insertions began after Jan Brandt, chief marketing officer, had an inspiration at the local Blockbuster video rental store. Says Brandt:

> After a very long work week, I went into my local Blockbuster to rent some movies to help relax over the weekend. They had sample boxes filled with popcorn, small candies and coupons that they were giving away if you rented 3 movies. I took the box into work that Monday and we set out to get our disks inserted into these boxes. We got over a 3% response on this effort (for context, most insert programs like that would yield in the tenths of a percent). The dramatic success of this effort led to the acceleration of the alternative marketing programs we became known for where you'd find disks popping up in what seemed like unlikely places.

See http://www.quora.com/How-many-different-AOL-promo-discs-were-distributed, accessed November 2012.

[25]According to various recollections, this cost AOL approximately $35 per customer (Steve Case, CEO), with the firm logging in a new customer every six seconds (Jan Brandt, chief marketing officer). When AOL 4.0 was launched in 1998 AOL bought all the CD-ROMs produced on the planet for several weeks (Reggie Fairchild, product manager, AOL 4.0), yielding eight million additional users in a year. See http://www.quora.com/AOL-History/How-much-did-it-cost-AOL-to-distribute-all-those-CDs-back-in-the-1990s, accessed November 2012.

As a further example, in the mid- to late 1990s some cable companies believed they did not understand Internet users' requirements, so they ceded these decisions initially to others, such as, for example, @home, a new entrant that positioned its service for cable firms. Eventually, @home merged with Excite to gain access to the perceived advantage of owning a portal, a decision that was regretted later by several cable firms. When the cooperation between cable firms and @home/Excite ended, it produced a large amount of recrimination, and the transition was not smooth for users.[26]

Although this experiment was not financially successful for @home, the surviving firms—cable companies, in this case—learned valuable lessons about how to structure their ISP services. First, certain useful "investments" were recreated, such as geographic caching of content, and, second, certain "mistakes" were avoided, such as advertising. Cable companies depended on subscription revenue for access services thereafter.[27]

Many local ISPs sought to differentiate themselves from one another by offering additional services, such as hosting, web development, network services, high-speed access, and tailored content.[28] These services were in addition to e-mail, newsgroup service, or easy access to portal content, online account management, customer service, technical support, Internet training, and file space.[29]

Once again, ISPs in the major urban areas behaved differently from those in less dense settings. Table 8.1 shows the result of a survey for national, urban, and rural ISPs in 1998.[30] Virtually every ISP offered basic service, but the extras were far less common among small and/or rural ISPs. In general, a much lower percentage offered hosting, web development, network services, or high-speed services. In short, ISPs in the urban areas were more differentiated than those in the rural areas, and the larger ISPs also tended to be located more often in urban areas.

What lay behind the differences? ISPs approached the new opportunities with employees and entrepreneurs who possessed different skills, experiences, and commercial priorities. Some ISPs already had successful

[26] See, e.g., Rosston (2009).

[27] See, e.g., Rosston (2009) for an analysis of the changing views of cable firms about the source of value from controlling or not controlling a portal and ISP.

[28] See, e.g., Greenstein (2000b), O' Donnell (2001).

[29] For a review of these and related technologies, see Haigh (2008).

[30] These are a combination of the tables in Greenstein (2000a, 2000b).

PC businesses, or businesses offering network support services. Many of these investments could commit ISPs to offer a particular array of services, even before the full size of market demand was realized or the value of new commercial opportunities was fully known.

The ISPs also found themselves in widely different local circumstances. Areas differed in the number of ISPs, which led to very different local competitive pressure. In areas with multiple ISPs, not all came from the same background, leading to differences in approaches to differentiation. ISPs also could adopt different approaches to fostering long-term relationships with their customers, as well as anticipate different approaches to gaining revenue from such customers. Demand could vary over localities also, as some locales had many technical users and others had few.

An ISP's approach emerged out of combinations of these firm-specific and location-specific factors. Some ISPs would focus on specific customers, offering a cluster of services for that user base—for example, offering network support for small businesses in the area. Others tailored their services to a wide group of local customers—for example, offering hosting services to any household or small business that wanted to operate a web page. Some specialized in handholding activities, helping nontechnical customers walk through various stages of getting online, sending e-mail, and surfing. National ISPs most closely resembled their counterparts among urban local ISPs. As table 8.1 shows, national ISPs too had a higher propensity to offer additional services.

Most attempts at differentiation did not succeed at raising the price. At most, these helped attract customers or preserve market share. For example, small ISPs were largely unable to raise prices as the Internet improved. Many ISPs improved their own service in myriad ways, as did many complementors.[31] Yet such improvements were part of the standard package rather than vehicles for premium pricing.

An excellent illustration was the upgrade from 28K to 56K modems, which began in 1997 and accelerated in 1998. There were two price levels for a short time—with a higher premium for faster modem service, but by

[31]Most ISPs had adopted technologies to enable web pages to upload more quickly, either by improving dynamic web page allocation through caching or by making arrangements with Akamai for its service. In addition, browsers got better, and so did Apache, so every user's experience improved.

TABLE 8.1. Service lines of ISPs, 1998

Category definition	*Most common service offered*	*ISPs, weighted by served territory*	*All ISPs*	*Small ISPs*	*Rural Small ISPs*
Providing and servicing access through different channels	28.8, 56K, isdn, web TV, wireless access, T1, T3, DSL, frame relay, e-mail, domain registration, new groups, real audio, ftp, quake server, IRC, chat, video conferencing, cybersitter TM	28,967 (100%)	3,816 (100%)	2,089 (100%)	325 (100%)
Networking service and maintenance	Networking, intranet development, WAN, co-location server, network design, LAN equipment, network support, network service, disaster recovery, backup, database services, novell netware, SQL server	8,334 (28.8%)	789 (20.6%)	440 (21.1%)	36 (11.0%)
Website hosting	Web hosting, secure hosting, commercial site hosting, virtual ftp server, personal web space, web statistics, BBS access, catalog hosting	8,188 (28.2%)	792 (20.7%)	460 (22.0%)	45 (13.8%)
Web page development and servicing	Web consulting, active server, web design, java, perl, vrml, front page, secure server, firewalls, web business solutions, cybercash, shopping cart, Internet marketing, online marketing, electronic billing, database integration	13,809 (47.7%)	1,385 (36.3%)	757 (36.2%)	76 (23.3%)
High speed access	T3, DSL, xDSL, OC3, OC12, Access rate > 1,056K	15,846 (54.7%)	1,059 (27.8%)	514 (24.6%)	39 (12.0%)

Note: The category definitions give the most general aggregation of services, while the most common service offered provides self-described phrases for ISPs. Column 3 weights an ISP by the number of area codes it serves, while column 4 treats each ISP as a single unit. Column 5 examines a subset of column 4, small ISPs, defined as any ISP covering less than five area codes. Column 6 examines a subset of Column 5, small ISPs outside of the over five hundred urban counties, where urban is defined as a MSA by the 1990 census.

late 1998, the twenty-dollar price umbrella prevailed once again.[32] In other words, the primary gatekeepers for access did not capture a higher fraction of the value from improvements in the network's many components.

Although most ISPs could not raise prices, few saw any reason to lower them—especially if users were reluctant to give up e-mail addresses or other services to which they had grown accustomed. Accordingly, many ISPs adopted policies refusing to forward e-mail for former customers, as a way to make users reluctant to switch ISPs.

What else kept prices from falling more rapidly?[33] By continuing to invest in content inside their walled gardens ISPs such as AOL recognized the value of those gardens for their customers and endeavored to keep them from migrating elsewhere. Switching e-mail was costly, and AOL also invested in and facilitated community development with services, such as instant messaging and chat rooms. Once ensconced in these applications with a familiar group of friends, users were reluctant to leave AOL and comfortable communities. In addition, instant message users did not want to recreate their buddy lists for another platform.

Growing a New Source of Revenue

No single administrative agency could possibly have built and managed the commercial network that emerged after 1995. The shape, speed, growth, and use of the commercial Internet after 1995 were not foreseen within the government circles responsible for its birth. That was not due to any oversight or error, and it occurred in spite of comparatively benign motives from the overseers, as well as abundant advice from the best technical experts in the world.

The commercial network grew rapidly. Access fees generated most of the revenue during the first decade of the commercial Internet. The typical household spent more than three-quarters of its time online at free or

[32] Stranger and Greenstein (2006).

[33] The declines each year were no more (on average) than a dollar reduction for a monthly contract. See Stranger and Greenstein (2006) for estimates of price trends for non-AOL services. The Bureau of Labor Statistics began a price series for Internet access after December 1997. The series is dominated by the pricing of the largest firms (AOL in particular), so this series shows less than a half of 1 percent decline in prices between its inception in December 1997 and December 2002.

TABLE 8.2. Adjusted revenue for access markets, 1998–2003 (millions of dollars)

Year	1998	1999	2000	2001	2002	2003
Dial-up	5,499	8,966	12,345	13,751	14,093	14,173
DSL		228	1,245	2,822	4,316	6,954
Cable modem	138	274	903	2,600	4,117	7,372

ad-supported sites, devoting most of its Internet budget to access fees, not subscription fees.[34]

The US Census Bureau began tracking ISP revenue as a separate revenue category in 1998. Table 8.2 provides an adjusted summary of these reports.[35] Several sources suggest that between 60 percent and 75 percent of the revenue in table 8.2 came from households, depending on the year and access mode. The remainder comes from business users of Internet access. The growth in revenues in table 8.2—from $5.5 billion in 1998 to over $28 billion in 2003—is astonishing for an entirely new market, especially one that did not start growing quickly until after 1995. Broadband revenues constitute approximately half the total revenue over the eight years, beginning with less than 6 percent in 1999 and growing to half of the total revenue in 2003.

To say it simply, the adaptation and circular conundrum were overcome with remarkable alacrity. A range of foresighted and complementary investments came together in a short period—the newly privatized Internet and its mesh, the newly privatized browser in commercial form, a few intrepid websites and portals to organize the experience, and commercial ISPs with a digestible form of the service that generated revenue. Together these activities added up to a commercial system nobody could have forecast. The Internet became attractive to more than just technical users. Most astonishing of all, it became available almost everywhere in the United States.

[34]See Goldfarb (2004).
[35]For all the details see the appendix to Greenstein and McDevitt (2012).