Vertical Integration," forthcoming *International Journal of Industrial Organization* (1999); Nicholas Economides and Steven Salop, Competition and Integration Among Complements, and Network Market Structure, 40 J. *Industrial Economics* 105 (1992); Randall Heeb, Innovation and Vertical Integration in Complementary Software Markets (Ph. D. dissertation, University of Chicago, 1999).

14. The terms were originally defined supra note 1.

15. The papers cited supra note 13 make this point with respect to their models as well.

11

Empirical Evidence on Advanced Services at Commercial Internet Access Providers

Shane Greenstein

This study analyzes the service offerings of Internet Service Providers (ISPs), the leading commercial suppliers of Internet access in the United States. It presents data on the services of 3816 ISPs in the summer of 1998. By this time, the Internet access industry had undergone its first wave of entry and many ISPs had begun to offer services other than basic access. This chapter develops an Internet access industry product code which classifies these services. Significant heterogeneity across ISPs is found in the propensity to offer these services, a pattern with a large/small and an urban/rural difference. Like in other (telecommunications) services, rural ISPs provide less choice than urban. The question therefore is if the Internet poses universal service problems similar to telephony.

1 Introduction

This study sheds light on the services at 3816 Internet Service Providers (ISPs) in the summer of 1998. Commercial providers account for the vast majority of Internet access in the United States. While revenues, estimated between $3 and $5 billion in 1997 (Maloff 1997), are relatively small for the communication and computing industry, they are rather large for a four year old industry.

More specifically, this chapter investigates the propensity of an ISP to offer services other than routine and basic access. These services represent the response of private firms to new commercial opportunities and technical bottlenecks. I investigate four types of services: frontier access, networking, hosting and web design services. Because no government agency, such as the Census or the Bureau of Labor Statistics, has yet

completely categorized these services, this study is the first economic analysis to develop and employ a novel Internet access product code. This classification leads to the first-ever documentation of two important patterns: first, there is no uniformity in ISPs' experiments with nonbasic access services; second, the propensity to offer new services shows both a notable urban/rural difference and large/small firm difference. These findings are, by themselves, important to on-going policy discussions about the development of the Internet infrastructure market.

Services from ISPs are an excellent example of how Internet technology had to be packaged in order to provide value to commercial users. When the Internet first commercialized it was relatively mature in some applications, such as e-mail and file transfers, and weak in others, such as commercial infrastructure and software applications for business use. This was due to the fact that complementary Internet technology markets developed among technically sophisticated users before migrating to a broad commercial user base, a typical pattern for new information technology (Bresnahan and Greenstein 1999). The invention of the World Wide Web in the early 1990s further stretched the possibilities for potential applications, exacerbating the gap between the technical frontier and the potential needs of the less technically sophisticated user.

Many ISPs pursued distinct approaches to developing commercial opportunities, which industry commentators labeled "different business models." This variety arose because, unlike the building of every other major communications network, Internet infrastructure was built in an extremely decentralized market environment. Aside from the loosely coordinated use of a few *de facto* standards (e.g., World Wide Web), government mandates after commercialization were fairly minimal. In every major urban area in the U.S. hundreds of ISPs built, operated, and delivered Internet applications, tailoring their network offerings to local market conditions and entrepreneurial hunches about growing demand. Not surprisingly, in the first four years after the commercialization of the Internet, the products changed frequently, many firms changed strategies, and the market did not retain a constant definition. This study provides a benchmark for understanding and measuring the observed variety among ISPs.

2 The Internet Access Business after Commercialization

Internet technology is not a single invention, diffusing across time and space without changing form. Instead, it is a suite of communication technologies, protocols and standards for networking between computers. This suite is not valuable by itself. It obtains economic value in combination with complementary invention, investment and equipment. How did Internet technology arise and how did these origins influence the commercialization of the technology?

The Origins of Internet Technology

By the time of commercialization, Internet technology was a collection of (largely) nonproprietary *de facto* standards for the development of communications between computers. These arose out of DARPA (Defense Advanced Research Projects Agency) experiments aimed at developing communications capabilities using packet switch technology. In 1969 DARPA began the first contracts for ARPANET, which involved a few dozen nodes. The first email message arrived in 1972. After a decade of use the protocols that would become TCP/IP were established and in regular use. By 1984 the domain name system was established and the term Internet was used to describe the system. In the early 1980s DOD began to require the use of TCP/IP in all Unix-based systems which were in widespread use among academic research centers.

In 1986 oversight for the backbone moved to the NSF, leading to a dismantling of ARPANET, and the establishment of a series of regional networks. The NSF pursued policies to encourage use in a broad research and academic community, subsidizing access to the Internet at research centers outside of universities and at nonresearch universities. The NSF policies had the intended effect of training many network administrators, students and users in the basics of TCP/IP technology. The NSF also sponsored development of changes to TCP/IP that enabled it to apply to more varied uses. Thus, this period saw the development of a variety of disparate technologies, most of which embodied nonproprietary standards, reflecting the shareware, research or academic culture in which they were born. Most of these would soon become necessary for the provision of basic access.

The unanticipated invention of the World Wide Web associated a new set of capabilities, display of nontextual information, with Internet technology. This was first invented in 1989 for the purpose of sending scientific pictures between physicists (though some alternatives were also under experimental use at the time). By the time the Internet was commercialized, a new set of experiments with browsers at the University of Illinois had developed the basis for Mosaic, a browser using web technology, something which made the whole suite of Web technologies easier to access. Mosaic was widely circulated as shareware in 1993–94 and quickly became a *de facto* standard, exposing virtually the entire academic community to the joy of sending pictures. The commercial browsers that eventually came to dominate nontechnical use, Netscape and Internet Explorer, sprang from these technical beginnings.

By 1995 there was an economic opportunity to create value by translating the basic pieces of Internet technology into a reliable and dependable standardized service for nontechnical users. This involved building access for business and home users. It also involved solving problems associated with customizing TCP/IP to networks in many different locations running many distinct applications. The primary open issues were commercial, not technical. Was this commercial opportunity fleeting or sustainable? What business model would most profitably provide Internet access, content and other services to users outside the academic or research environment? What services would users be willing to pay for and which services could developers provide at low cost?

Adaptation Activity in the Internet Access Market after Commercialization

As it turned out, market-based transactions quickly became the dominant form for delivery of on-line access. Commercial ISPs developed a business of providing Internet access for a fee. Access took one of several different forms: dial-up to a local number (or a toll free number) at different speeds, or direct access to a business's server using one of several high-speed access technologies. Within three years the commercial providers almost entirely supplanted their academic parents. By the spring of 1998 there were scores of national networks covering a wide variety of dial-up and direct access. There were also thousands of regional and local providers of Internet access that served as the links between end-users

and the Internet back-bone (see Downes and Greenstein 1999 for detail). As of 1998, less than 10% of U.S. households and virtually no business get Internet access from university-sponsored ISPs (Clemente 1998).

In retrospect, several economic factors shaped this entry. Technology did not serve as a barrier to entry, nor were there prohibitive costs to hiring mainstream programming talent. Providing basic access required a modem farm, one or more servers to handle registering and other traffic functions, and a connection to the Internet backbone.[1] Some familiarity with the nonproprietary standards of the web was required, but not difficult to obtain. Because so many students had used the technology in school, and because the standards were nonproprietary, anyone with some experience could use them or learn them quickly. As a result, a simple dial-up service was quite cheap to operate and a web page was quite easy to develop (Kalakota and Whinston 1997).

The amateurs of 1995 soon learned that cheap and easy entry did not necessarily translate into a profitable on-going enterprise. The major players from related markets who opened large access services, such as AT&T, also learned that the basic access market had small margins. By 1998 basic access was not generally regarded as a very lucrative part of the ISP commercial market in virtually any location.

By 1998 different ISPs had chosen distinct approaches to developing access markets, offering different combination of services. Why did this variance arise? Answering these questions provide a window on the factors shaping adaptation activity in Internet Technologies.

3 Varieties of Business Models in Technology-Intensive Markets

Standard economic analysis offers a number of explanations for why different firms pursue different strategies for adapting to the diffusion of a general purpose technology. As emphasized in Bresnahan, Stern, and Trajtenberg 1997, one way to frame such an empirical investigation is to view an ISP's choices as an attempt to differentiate from common competitors. Firms may try to push technical frontiers, develop local or national brand names, combine recent technical advances with less technical businesses and so on. Such differentiation may arise as a response to firm-specific or user-specific assets, and these returns may be temporary if competitors eventually learn to provide close substitutes.

The 1998 Internet access industry can be understood in these terms, though the framework also needs modification to account for important features of Internet infrastructure markets. Entry into many locations had extensively developed the "basic access" market, the first and most obvious adaptation of Internet technologies to commercial use. Due to this extensive entry, the private returns to basic access services in most locations had almost entirely been competed away by 1998. Thus, supernormal private returns, if they existed at all, existed in differentiating from basic access.

From the viewpoint inside a firm, what an ISP does in a particular market situation is an strategic question. In contrast, what all firms do across the country is an empirical economic question. This study uses the industry discussion of strategy as a basis for understanding market wide behavior. That is, industry trade publication distinguish between two types of activities other than basic access.

- *Offering technically difficult access:* High-bandwidth applications present many technical difficulties which challenge the skills and capital constraints of many ISPs. The slow diffusion of commercially viable high-speed access markets is widely regarded as a major bottleneck to the development of the next generation of Internet technologies. Accordingly, this type of commercial offering has generated much policy interest (Esbin 1998).
- Offering services that are complementary to basic access: Providing additional services became essential for retaining or attracting a customer base. Many ISPs instead tried to develop additional services, such as web-hosting, web-design services and network maintenance for businesses. Any of these were quite costly, as they had to be properly assembled, maintained, and marketed. Because many of these services push the boundaries of existing telecommunications and computing market definitions, these too have generated much policy interest (Werbach 1997).

Which factors determined the provision of nonbasic access services? In their theory of general purpose technologies and coinvention, Bresnahan and Trajtenberg (1995) place emphasis on the dispersion of factors that change incentives at different locations, between firms, and over time. That is, many firms and locations face the same secular technological trends, hence they share similar technical factors. Suitably altered to this study's situation, this framework predicts that *differences* across firms at any point in time (or over time) arise when decision makers face different incentives arising from differences in demand conditions, differences in the quality of local infrastructure, differences in the thickness of labor markets for talent, or differences in the quality of firm assets. These create a variety of economic incentives for adapting Internet infrastructure to new uses and applications. Greenstein (1999) investigates statistical models for understanding the determinants of these observed differences.

Generating the Original Sample

To characterize the offering of service in a quantitative way, I and some research assistants examined the business lines of 3816 Internet service providers in the United States who advertise on *thelist* (see Greenstein 1999, appendix I and II for details). This site, maintained by Meckler Media, provides the opportunity for both large and small ISPs to advertise their services. ISPs fill out a questionnaire where the answers are partially formatted, then the answers are displayed in a way that allows users to compare different ISP services.

This group of 3816 ISPs will be called the "original sample." This study also contains additional information for a subset of them labeled the "analysis sample." This group has 2089 ISPs. Its construction will be described in detail below. Virtually every firm in the original and analysis samples provides some amount of dial-up or direct access and basic functionality, such as email accounts, shell accounts, IP addresses, new links, FTP, and Telnet capabilities.

From comparison with other sources, such as *Boardwatch, thedirectory* and the National Telephone Cooperative Association directory on Internet Services in rural areas (NTCA 1998), it appears that these 3816 ISPs are not a comprehensive census of every ISP in the country. The Downes and Greenstein (1999) sample of the ISP market in the spring of 1998, which is constructed primarily from information culled off *thedirectory*, found over 6100 ISPs in the United States. These 3816 seem to under-represent ISPs in small towns (e.g., where advertising on the web is not necessary) and quasi-public ISPs (e.g., rural telephone companies[2]). In addition, this sample does not examine firms who offer nonbasic services but who do not offer basic access. That said, it does contain many observations from small firms, from ISPs in rural areas and from virtually all the mainstream ISPs from whom the vast majority of Internet users in the United States get their access.

Classifying the Services of ISPs

The first goal is to classify the activities of Internet access firms. No product code exists for this industry, as it has grown faster than government statistical agencies can classify it. Based on trade literature and magazines, I grouped services into five broad categories: basic access, frontier access, networking, hosting, and web page design (see appendix II for the product code).

- *Basic access* constitutes any service slower than and including a T–1 line. Many of the technologies inherited from the pre-commercial days were classified as complementary to basic access, not as a new service.
- *Frontier access* includes any access faster than a T–1 line, which is becoming the norm for high-speed access to a business user. It also includes ISPs which offer direct access for resale to other ISPs or data-carriers; it also includes ISPs who offer parts of their own "backbone" as a resale to others.[3]
- *Networking* involves activities associated with enabling Internet technology at a user's location. All ISPs do a minimal amount of this as part of their basic service in establishing connectivity. However, an extensive array of these services, such as regular maintenance, assessment of facilities, emergency repair, and so on, are often essential to keeping and retaining business customers. Note, as well, that some of these experimental services could have been in existence prior to the diffusion of Internet access; it is their offering by an Internet access firm that makes them a source of differentiation from other ISPs.
- *Hosting* is typically geared toward a business customer, especially those establishing virtual retailing sites. This requires the ISP to store and maintain information for its access customers on the ISP's servers. Again, all ISPs do a minimal amount of hosting as part of basic service, even for residential customers (e.g., for email). However, some ISPs differentiate themselves by making a large business of providing an extensive array of hosting services, including credit-card processing, site-analysis tools, and so on.
- *Web Design* may be geared toward either the home or business user. Again, many ISPs offer some passive assistance or help pages on web page design and access. However, some offer additional extensive consulting services, design custom sites for their users, provide services associated with design tools and web development programs. Most charge fees for this additional service.

Other services were put into four other groups: traditional computing services (e.g., PC sales and service), traditional telecommunications (e.g., cellular phone sales and service), consulting, and miscellaneous services (e.g., copying, cafes and photography). While in practice these last four were less common, the nonaccess lines of business of ISPs will be useful. For the most part, if an ISP advertises this business service, this was this firm's primary business before the firm became an ISP.

Descriptions of each ISP's services on *thelist* was classified into standard "phrases" which are then mapped to particular services at particular ISPs. In other words, an ISP offers networking services if that ISP uses one of the "phrases" which corresponds to networking activity. Similar exercise followed for hosting, web design, frontier access and so on. An ISP could be in more than one service. Table 11.1 lists the most common phrases for each line of business. (The entire list of phrases and the correspondence table are available from the author on request. See Greenstein 1999, appendix II for the product code.) In general, these methods should *undercount* the offering of any particular service line since many phrases were uninformative. In other words, this method will only record a service line if the ISP clearly states it as such.[4] In addition, the lines between different services are often, but not always, sharp. This warrants a cautious interpretative approach, because ambiguities in definitions naturally arise.

By definition, every ISP has at least one useful phrase indicating activity in the access business. On average, an ISP had 8.6 useful phrases (standard deviation of 4.6, maximum of 40). The main statistical findings from applying the classification scheme are listed in table 11.1 for three different samples, including the original sample. These findings are also illustrated by figures 11.1a and 11.1b.

A First Look at the Service Lines of ISPs

Of the 3816 firms in the original sample, 2295 (60.1%) have at least one line of business other than basic dial-up or direct Internet access. Table 11.1 shows that 1059 provide high speed access, 789 networking, 792 web hosting, 1385 web page design. There is some overlap (shown in figures 11.1a and 11.1b): 1869 do at least one of either networking, hosting or web design; 984 do only one of these three; 105 do all three and frontier access. The analysis sample has similar percentages. For such a cautious method, this reveals quite a lot of experimentation with nonaccess services by firms in the access business.[5]

Table 11.1
Product Lines of ISPs

Category definition	Most common phrases in category	Weighted by service territory*	Original sample	Analysis sample**
Providing and servicing access through different channels	28.8, 56k, isdn, web TV, wireless access, T1, T3, DSL, frame relay, e-mail, domain registration, news groups, real audio, ftp, quake server, IRC, chat, video conferencing, cybersitter TM	28967 (100%)	3816 (100%)	2089 (100%) Rural ISPs *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	8334 (28.8%)	789 (20.6%)	440 (21.1%) Rural ISPs *(11%)*
Web site hosting	Web hosting, secure hosting, commercial site hosting, virtual ftp server, personal web space, web statistics, BBS access, catalog hosting	8188 (28.2%)	792 (20.7%)	460 (22.0%) Rural ISPs *(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	13809 (47.7%)	1385 (36.3%)	757 (36.2%) Rural ISPs *(23.3%)*
High speed access	T3, DSL, xDSL, OC3, OC12, Acccess rate>1056k	15846 (54.7%)	1059 (27.8%)	514 (24.6%) Rural ISPs *(12.0%)*

Notes: *Unit of observation is ISP-Area codes, as found in *thelist*. For example, if an ISP offers local dial-up service in 29 area codes, it will be 29 observations. If that same ISP offers high speed access then it will count as 29 cases of high speed access.

**Unit of observation is an ISP in small number of territories. See text for precise definition. Top number is for all 2089 ISPs in analysis sample. *Italicized* percentage is for the 325 ISPs found primarily in rural areas.

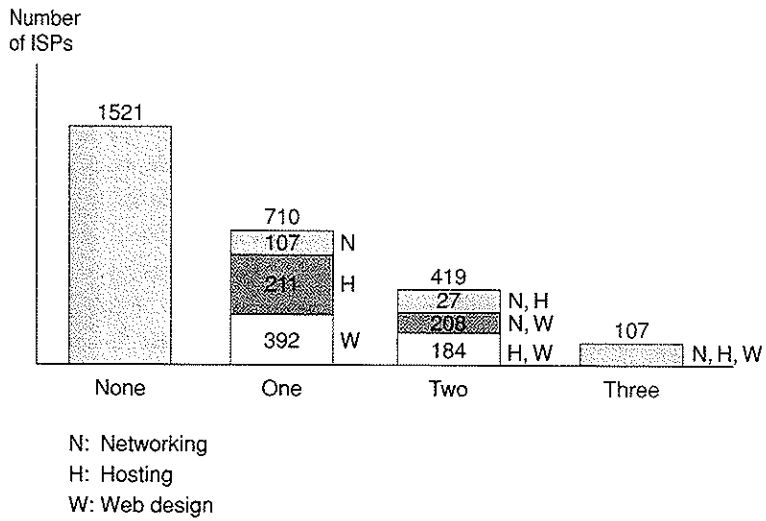

Figure 11.1a
Experiments with new services by ISPs without frontier access technology

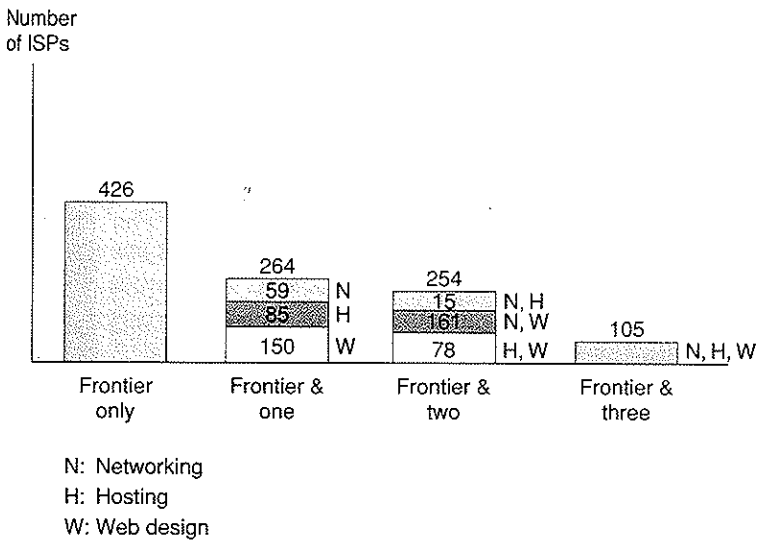

Figure 11.1b
Experiments with new services by ISPs with frontier access technology

The largest firms—defined as present in 25 or more area codes—experiment at slightly higher rates: 159 of 197 firms (in this sub-sample) are in either networking, hosting or web design—60 do only one, 18 do all four. 115 provide high speed access, 59 networking, 63 web hosting, 94 web page design. That is a higher rate than the whole sample, but consistent with the hypothesis that urban areas (where large firms are disproportionately located) tend to receive higher rates of experimentation from their ISPs. This hypothesis receives further attention below in the analysis sample.

The above indicates that using the ISPs as the unit of observation may provide a partly distorted view of the geographic diffusion of new services. To develop the point further, table 11.1 lists another column which weights experimentation—admittedly, coarsely—for geographic dispersion. The product line is weighted by the number of area codes in which the ISP provides service. Since this is the only data available about geographic dispersion for all 3816 ISPs in *thelist*, this is the most one can do. This weighting is coarse because not all area codes are equal in square miles, nor population.[6]

In the original sample, ISPs are in 7.6 area codes on average. There were 28,967 "ISP-Area Codes." Of these 17,343 (77.2%) have at least one additional line of business other than dial-up Internet access or routine direct access, higher than found in the un-weighted sample. Even emphasizing how cautious these methods are, this second way of representing the data reveals quite a lot of experimentation in nonaccess business. Table 11.1 shows that, using 28,967 as denominator, 15,846 ISP-area codes provide high speed access, 8334 networking, 8188 web hosting, 13,809 web page design. In all cases, these are higher percentages than the original sample; in the case of high speed access, this is a much higher percentage. Because the firms in a larger number of regions tend to do more experimentation, this suggests that most users, especially those in urban areas where the national firms tend to locate, probably have access to some form of experimentation.[7]

These first results do not seem to be an artifact of survey bias. There is not enough evidence here to suggest something artificial about the relationship between the results and the effort it takes to fill out the survey form for *thelist*.[8]

The Relationship between Services

Table 11.2 examines the two different types of services other than basic access. Specializing in very high-speed Internet services is one type of service that distinguishes a firm from its competitors. As has been noted in many places (e.g., Kalakota and Whinston 1997), greater and greater speeds are harder to obtain (as a technical matter) and costly to reliably provide (as a commercial matter). In contrast, specializing in hosting, networking service or web design can also distinguish a firm from its competitors. Many of these services require trained personnel and may be difficult to do profitably. Hence, these two types of differentiation might be done by the same firms for commercial reasons, but there is no necessary technical reason for it.

In the original sample the fraction of firms in networking, hosting, and web design is higher among those in high speed access, but the relationship is not very strong. Table 11.2 shows that in the original sample, of the 1059 in high speed access, 59.8 percent (633) provided networking, hosting or web design. By comparison, of the 2757 not providing high-speed access, 44.8 percent (1236) did so. Similarly, of the 1869 providing networking, hosting or web design services, 33.9 percent (633) provided high speed access. Of the 1947 not providing networking, hosting and web design, 21.8 percent (426) provided high speed access.

For comparative purposes, table 11.2 also lists the same correspondence for the data in ISP-area codes and for the analysis sample. The same qualitative results remain. Comparison of different lines of business with each other is shown in figures 11.1a and 11.1b. These reinforce the point that different firms carry different nonbasic services. The determinants of these patterns are discussed further below.

Constructing the Analysis Sample

Additional data about local conditions is available for the analysis sample. However, this additional information comes at the cost of a reduced sample size and with a potential selectivity bias.

The analysis sample was constructed as follows: First, the original sample was restricted to 3300 ISPs in 20 or fewer area codes, as found in *thelist*. This isolates regionally dispersed decision makers. Second, the original sample was compared against a set of roughly 5400 ISPs in the Downes and Greenstein (1999) data set for ISPs, which were in five or

Table 11.2
Product Lines of ISPs

Original sample

		Network, Hosting & Web		
		Offers	Does not	Total
High speed offers frontier		633	426	1059
Access	Does not	1236	1521	2757
	Total	1869	1947	3816

Weighted by Service Territory

		Network, Hosting & Web		
		Offers	Does not	Total
High speed offers frontier		10822	5024	15846
Access	Does not	6521	6600	13121
	Total	17343	11624	28967

Analysis Sample

		Network, Hosting & Web		
		Offers	Does not	Total
High speed offers frontier		314	200	514
Access	Does not	736	839	1575
	Total	1050	1039	2089

fewer counties. This small geographic setting allows the accurate identification of the degree to which the ISP serves high ("urban") or low ("rural") density areas. The Downes and Greenstein (1999) dataset for small ISPs comes from 1998 spring/summer listings in *thedirectory*, another forum in which ISPs advertise. *Thedirectory* places emphasis on listing the local dial-up phone numbers for many ISPs,[9] which permits identification of the local points of presence (POPs) for ISPs, and, hence, the local geographic territories served by any ISP who offers dial-up service.[10] This is a much finer way to identify local service territories and local market conditions than using the area codes included in *thelist*.[11]

Third, an ISP was included in the analysis sample if the ISP listed the same domain name for the home page in both *thedirectory* and *thelist*. This strategy ensures 100% accurate observations. The emphasis on accuracy was deliberate. It was discovered that it is relatively common for several different firms to maintain similar company names and similar domain names, heightening potential confusion. In addition, many ISPs maintain several similar "home pages" with different domain addresses for a variety of reasons (e.g., tracking traffic from different sources, marketing under different organizational umbrellas, etc.). Since the dataset is large enough for the statistical purposes below (2089 observations) and there was no hope of getting a census of all ISPs, the benefits of absolute accuracy overwhelmed the potential benefits of a mildly larger sample which ran the risk of being inaccurate for a few firms.[12]

These 2089 ISPs are representative of small ISPs. Comparisons of the 2089 ISPs in the analysis sample with the roughly 5400 small ISPs in Downes and Greenstein 1999 showed little difference in the features of the service territories. In the analysis sample 83.5% of the ISPs are in urban counties, using the broadest definition of urban from the U.S. Census. In Downes and Greenstein, only 81.1% are in urban counties. Other than this slight difference, there is no qualitative difference in the average features of the territories covered by small ISPs in the two data sets. Moreover, the number of small ISPs found in each county in the Downes and Greenstein dataset and in the analysis sample correlates at .94, as one would expect if the 2089 ISPs in the analysis sample were nearly a random selection of small ISPs from across the country.[13] In sum, the two known biases in the analysis sample are the slight over-representation of urban areas, for which it is possible to control, and the

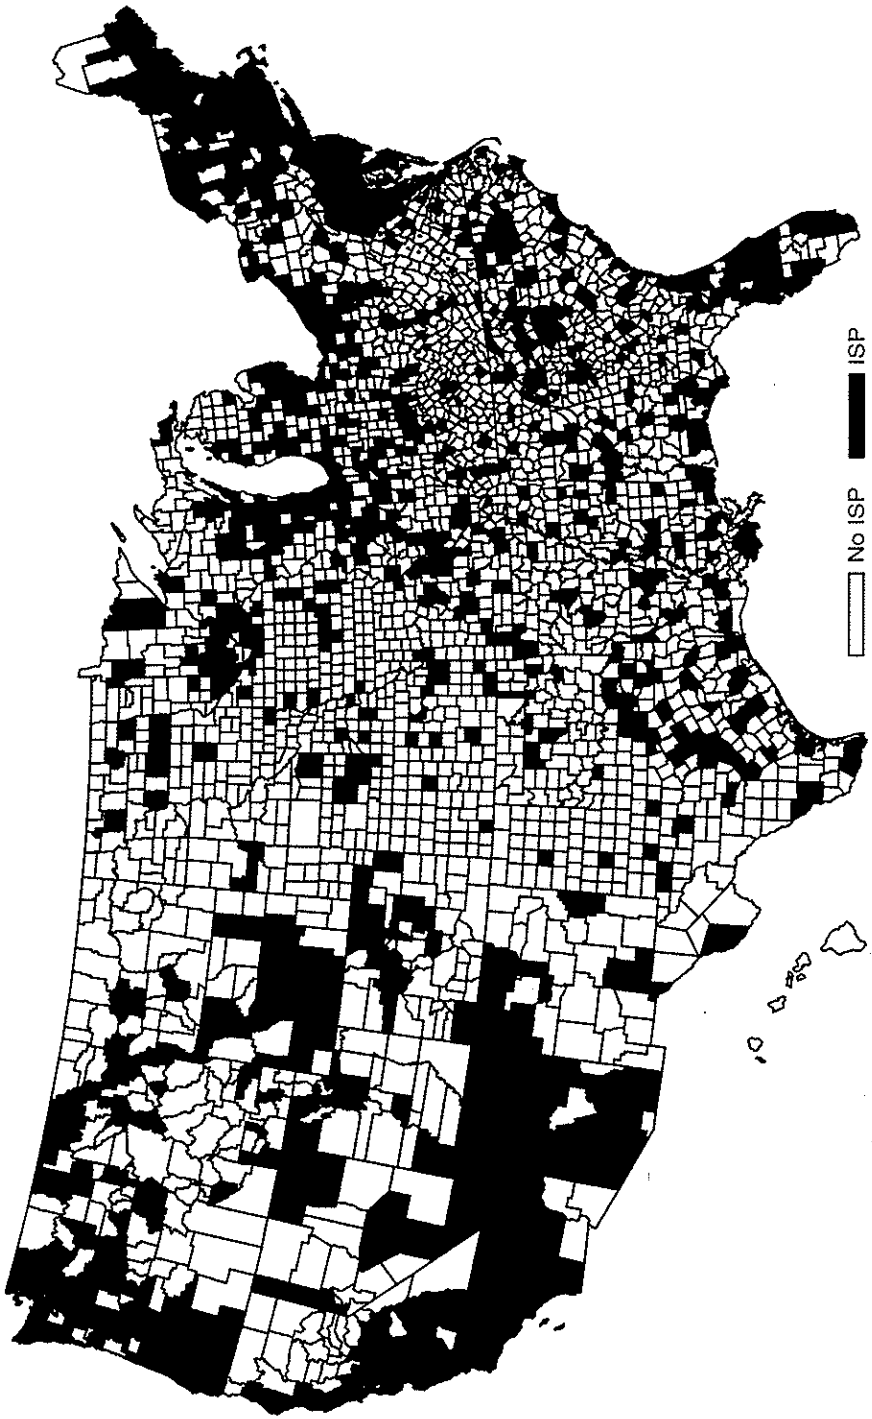

sample bias toward small firms, which I choose deliberately so I could identify local conditions.[14]

Figure 11.2 presents a map where a county is blackened if it contains at least one ISP from the analysis sample. There are 905 such counties represented in this sample, with representatives from virtually every urban area in the U.S. as well as several hundred rural counties.

A Second Look at the ServiceLines of ISPs

Tables 11.1 and 11.2 present comparable statistics for experimentation by ISPs in the analysis sample. These tables show patterns similar to those for the original sample. This evidence suggest that the analysis sample is not an unrepresentative sub-sample of the original sample.

The focus on small firms permits a close examination of differences in experimentation by urban and rural ISPs. The results are striking. In this sample of 2089, 1764 ISPs primarily serve urban areas.[15] Their propensities to offer services are slightly higher than for the whole original sample. Of those 1764 ISPs, 26.9% offer frontier access, 22.9% offer networking services, 23.5% offer hosting services, and 38.6% web design services.

The last column of table 11.1 shows the contrast with rural areas. Of the 325 ISPs primarily found in rural areas, 12.0% offer frontier access, 11.0% offer networking services, 13.8% offer hosting, and 23.3% offer web design services. The propensities are between 40% and 60% lower across every category. A simple test for difference of means between urban and rural ISPs strongly rejects the hypothesis that the average rate of experimentation is the same between the urban and rural samples of ISPs. This holds for every type of activity.

Implications for a Wider Debate

The above findings are novel, as no previous research has ever examined services at such a detailed service and geographic level. They raise important questions about the distribution of economic growth in Internet infrastructure markets. It is important to note, however, that the above findings are an unconditional comparison and say little about the determinants of

Figure 11.2
Presence of ISPs in analysis sample

outcomes. These findings do not control for urban/rural differences in population demographics, nor for urban rural differences in firm-specific traits. In other words, the geographic dispersion of the endogenous variable might be explained by geographic factors, but it might also be explained by the geographic distribution of firm-specific factors. Statistical analysis of this question may be found in Greenstein 1999.

This study's findings inform policy concern in the telecommunications industry (Werbach 1997, Esbin 1998, Weinberg 1999). They influence the creation and targeting of subsidies associated with new services, as proposed in the 1996 Telecommunications Act. If the lower propensity to find new services in low-density areas is due to an absence of local firms with appropriate skills, then policies might either induce ISPs to expand from high-density areas to low density areas (where they would not otherwise be), or it must induce incentives/vision/investments from ISPs who are already located in low-density areas (but who would otherwise choose to offer such services). If, on the other hand, the absence of new services in low-density areas is due to an absence of local demand for these services or the absence of local infrastructure, subsidies run the risk of not changing the propensity to experiment in such areas. Indeed, in that case, the subsidy can be very wasteful if it induces the offering of services which few want.

Second, and on a somewhat different note, much policy debate has been concerned with redefining the distinction between traditional telephone services and computing services (e.g., see Weinberg 1999, or Sidak and Spulber 1998, for a recent summary and critique). This distinction is a key premise of the 1996 Telecommunications Act, engendering FCC review of whether ISPs were special services, exempt from the access charges of telephone companies. Many authors have correctly noted that ISPs have benefitted from this exemption. It is a matter of debate about whether this exemption is welfare enhancing or decreasing.

This study has several comments to add to this extensive debate. By 1998 many ISPs were pursuing business models with only a mild relationship to the regulatory boxes and previous business lines (This study had to create its own new product code for precisely this reason.). This fact alone raises the question about the wisdom of employing these legacy regulatory categories to their behavior. It is, therefore, further evidence in favor of the arguments for constructing a new and possibly *sui generis* regulatory approach to ISPs. That said, if the regulatory and political process insists on trying to fit ISPs into one side or another of the historical line between telecommunications and computing, this study's approach adds one reason to the case for some forbearance and one to the case against. For forbearance: these experiments by ISPs may become a key market mechanism for developing the complementary Internet services that translate technical advance into economic value. As yet, it is unclear what shape these experiments will take next year and where these all will lead. It is in society's interest to have these experiments develop and in society's interest to let them have sufficient time to generate potential information spillovers. For settling the issue soon: many ISPs are developing their business models around a particular cost structure for a key input. It is in society's interest to have the cost of that input incorporated into the ISP industry's investments and other strategic commitments, then the distortions can be minimized if those costs are announced sooner instead of later.

4 Conclusion

Many technology enthusiasts have been waiting for the online revolution for a long time, welcoming the possibilities for new businesses, new services and new types of communications. Now that it is here, a commercialized Internet may not be precisely what they had in mind. Some locations have access to the latest technology from commercial firms and some do not, creating the potential for a digital divide. The economic benefits associated with new frontier technologies are diffuse, uneven and uncertain. Business use of the Internet is difficult and adaptation is time-consuming. Many new services do not employ frontier technology at all. Indeed, much commercialization involves bending frontier technology to the needs of commercial users, a process that often involves many non-technical issues.

The details of exploratory activity are inherently messy. In this case, ISPs customize Internet technologies to the unique needs of users and their organizations, solving problems as they arise, tailoring general solutions to idiosyncratic circumstances and their particular commercial strengths. Sometimes ISPs call this activity consulting, and charge for it separately, sometimes it is included as normal business practices. In

either case, it involves the translation of general knowledge about Internet technologies into specific applications which yield economic benefits to end-users. In all cases differences between their offering and their nearest competitor raise returns to innovative activity, inducing a variety of services from different ISPs.

Viewing the Internet access market in this way helps us to understand the explosive events just after the commercialization of the Internet. The technology underlying the Internet incubated in research laboratories but today's commercial industry has propelled it into common use. The economic value of ISP services is largely determined by the value commercial users place on it. This framework helps explain why the incubation of Internet technology in an academic setting lead to a lengthy set of adaptive activities in a nonacademic setting. These adaptations are hard to do, as they reflect ISP-specific capabilities and entrepreneurial guesses about the appropriate services to offer and about location-specific demands for particular services.

Notes

I would like to thank the Consortium for Research on Telecommunication Policy at Northwestern University for Financial Assistance. Thanks to Oded Bizan, Tim Bresnahan, Barbara Dooley, Tom Downes, Mike Mazzeo, Nancy Rose and Dan Spulber for useful conversations and to seminar audiences at the Federal Research Board Research Division in Washington DC, UC Berkeley, Northwestern and the conference on Competition and Innovation in the Personal Computer Industry. Angelique Augereau and Chris Forman provided outstanding research assistance. All errors are mine.

1. For example, see the description in Kalakota and Whinston [1996], Leida [1997], the accumulated discussion on <www.amazing.com/Internet/faq.txt>, or Kolstad [1998] at <www.bsdi.com>.

2. NTCA (1998) shows hundreds of rural telephone companies provide basic Internet services to their local areas, but it does not specify the extent of those services (for a further survey see, for example, Garcia and Gorenflo (1997)). Some of these rural telephone companies do advertise their services in either *the-list* or *thedirectory*, but a substantial fraction (> 50%) do not.

3. Speed is the sole dimension for differentiating between frontier and basic access. This is a practical choice. There are a number of other access technologies just now becoming viable, such as wireless access, which are slow but technically difficult. Only a small number of firms in this data are offering these services and these are coincident with offering high speed access.

4. The approach depended on the ISP describing in concrete terms the businesses they offer. For example, no additional line of business was assigned to an ISP who advertised "Call for details" or "We are a friendly firm." The vast majority of unused phrases were idiosyncratic phrases which only appeared once with one firm, defying general characterization. There were 1105 such phrases (and 6,795 unique useful phrases), which occurred 1406 times (out of 35,436 total phrases). In other words, most of the unused phrases occurred only once and described attributes of the firms which had nothing to do with their lines of business (e.g., HQ phone number, contact information or marketing slogans). The most common unused phrase was "etc."

5. One of the most difficult phrases to classify was general "consulting"—i.e., consulting that did not refer to a specific activity. Of all these vague consulting cases, all but 12 arose in the 1836 firms who provide networking, hosting and web design. Hence, the vast majority of consulting activity is accounted for by the present classification methods as one of these three complementary activities, networking, hosting and web-design.

6. Though there is, roughly speaking, a maximum limit on the total population associated with any given area code, this maximum only binds in a few locations. In general, therefore, area codes are not determined in such a way as to result in anything other than crudely similar population sizes and geographic regions.

7. However, this is only a hint and not a concrete conclusion. Without a complete census of new services at all ISPs, it is not possible to estimate precisely how much of the U.S. population has easy access to local provision of these services in the same way that Downes and Greenstein (1999) estimate the percentage of the population with access to basic Internet services.

8. Extreme geographic firm size (i.e., the total number of area codes in which the ISP offers service) is a good measure of a survey bias because the ISPs must expend effort to indicate the extent of their geographic coverage. If the number of phrases was low owing to ISP impatience with the survey format, one would expect a strong relationship between firm size and the number of phrases. Since the correlation is positive but small, which is plausible for many reasons having nothing to do with survey bias, I conclude that the number of lines of business does not arise as an artifact of ISP impatience with the survey or other forms of laziness by the ISP.

9. The other source of data for Downes and Greenstein (1999) is the *Boardwatch* backbone list, which concentrates mostly on national ISPs.

10. This is an artifact of the U.S. local telephone system, which tends to charge telephone calls by distance. Hence, the location of a local phone number from an ISP is an excellent indicator of the local geographic territory covered by the ISP. See Downes and Greenstein (1999) for further detail.

11. In some dense urban counties, the number of area codes exceeds the number of counties, but for most of the country the number of counties vastly exceeds the number of area codes. There are over 3000 counties in the U.S. and less than 200 area codes.

12. Because two sets of company names are maintained by two completely unrelated lists, *thedirectory* and *thelist*, each of whom uses different abbreviations and possibly different domain names, many ambiguities arose. It is certainly the case that many of the 1300 firms from *thelist* which are not included in the analysis sample are, in fact, in the Downes and Greenstein (1999) data. However, verifying these matches was tedious and potentially subjective, rendering it almost infeasible.

13. The correlation between the ISPs per county in the two datasets is .94 when Downes and Greenstein (1999) only examine ISPs in five counties or less. The correlation is, not surprisingly, lower when we correlate the number of ISPs in the analysis sample per county with the entire Downes and Greenstein dataset, which includes all national and regional firms. In that case, the correlation is .82. This is because larger firms tend to disproportionately locate in urban areas.

14. As noted in Downes and Greenstein (1999), there is also a subtle empirical bias in any study of ISPs. All inferences in this sample are conditional on observing the ISP in the access business to begin with. We do not observe those who considered this business, but did not choose it.

15. Each county an ISP serves is designated urban or rural by the U.S. census. In the rare cases where an ISP serves a mix of urban and rural areas, if the majority of counties are urban, then an ISP is said to be urban.

References

Bresnahan, Timothy, and Shane Greenstein (1999). "Technological Competition and the Structure of the Computing Industry," *Journal of Industrial Economics* 47:1–40.

Bresnahan, Timothy, Scott Stern, and Manuel Trajtenberg (1997). "Market Segmentation and the Source of Rents from Innovation: Personal Computers in the late 1980s," *Rand Journal of Economics* 28:s17–s44.

Bresnahan, Timothy and Manuel Trajtenberg (1995). "General Purpose Technologies: Engines of Growth?" *Journal of Econometrics* 65:83–108.

Boardwatch (1997). *March/April Directory of Internet Service Providers*, Littleton, CO.

Clemente, Peter C. (1998). *The State of the Net: the New Frontier*, McGraw-Hill, New York.

Downes, Tom, and Greenstein, Shane (1999). "Do Commercial ISPs Provide Universal Access?" in *Competition, Regulation and Convergence: Current Trends in Telecommunications Policy Research*, ed. Sharon Gillett and Ingo Vogelsang. Lawrence Erlbaum Associates, pp. 195–212.

Esbin, Barbara (1998). *Internet over Cable, Defining the Future in Terms of the Past*, FCC, Office of Planning and Policy Working Paper 30, August.

Garcia, D. Linda, and Gorenflo, Neal (1997). "Best Practices for Rural Internet Deployment: The Implications for Universal Service Policy," Prepared for 1997 TPRC, Alexandria, Va.

Greenstein, Shane (1999). "Building and Deliverying the Virtual World: New Services for Commercial Internet Access," <www.kellogg.nwu.edu/faculty/greenstein/images/research.html>.

Kalakota, Ravi, and Whinston, Andrew (1996). *Frontiers of Electronic Commerce*, Addison-Wesley, Reading, Mass.

Kolstad, Rob (1998). "Becoming an ISP," <www.bsdi.com>. January.

Leida, Brett (1997). "A Cost Model of Internet Service Providers: Implications for Internet Telephony and Yield Management," mimeo, MIT, Departments of Electrical Engineering and computer Science and the Technology and Policy Program.

Maloff Group International, Inc. (1997). "1996–1997 Internet Access Providers Marketplace Analysis," Dexter, Mo., October.

NTCA (1998). *Membership Directory and Yellow Pages*, National Telephone Cooperative Association, Washington, D.C.

Sidak, Gregory, and Spulber, Daniel (1998). "Cyberjam: The Law and Economics of Internet Congestion of the Telephone Network," *Harvard Journal of Law and Public Policy* 21, no. 2.

Weinberg, Jonathan (1999). "The Internet and Telecommunications Services, Access Charges, Universal Service Mechanisms, and Other Flotsam of the Regulatory System," *Yale Journal on Regulation*, Spring.

Werbach, Kevin (1997). "Digital Tornado: The Internet and Telecommunications Policy," FCC, Office of Planning and Policy Working Paper 29, March.