Innovative Conduct in Computing and Internet Markets

By Shane Greenstein[1]

Handbook of the Economics of Technical Change

Edited by Hall/Rosenberg

October, 2009

Abstract

How has innovative and competitive behavior in computing and internet markets evolved over the past half century? In the first section of this review, I discuss these questions in light of six topics: the limited role for technology push; the diffusion of general purpose technologies; the organization of proprietary platforms; the presence of asymmetric innovation incentives; the importance of market-oriented learning; and the localization of economic activity. Despite dramatic changes in outcomes, in the predominant product markets, and in the identities of leading sellers, the conditions of market structure shape innovative conduct in firms from one year to the next and, to a large extent, from one decade to the next, in many of the same economic terms.

In the second section, I closely examine the U.S. commercial Internet experience in the 1990s. While the peculiar events that led to the invention of the commercial Internet explain some of the salient and unique features of the commercial experience, much innovative activity resembles conduct seen for many decades in computing markets. This analysis highlights three additional topics: the division of technical leadership; the rise of open organizational forms for coordinating platforms; and the extraordinary breadth of activity touched by the Internet. These three additional factors account for many of the novel aspects of innovative conduct in Internet market.

Keywords: Computer Hardware, Computer Software, Internet, Innovation, Invention, Technology, Market conduct, Commercialization, Diffusion.

JEL Codes: L1, L63, L86, O3**.**

[1] Elinor and Wendell Hobbs Professor, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL. 60208. I thank Tim Bresnahan, Chris Forman, Avi Goldfarb, Bronwyn Hall, Rebecca Henderson, Jeff Prince, Nate Rosenberg, Alicia Shems, Kristina Steffenson McElheran, and Scott Stern for comments. The Searle Foundation provided financial support. I am responsible for all errors.

I. Introduction and Overview.

How has innovative and competitive behavior in computing and internet markets evolved over the past half century? This broad question does not and cannot have a simple answer for at least two reasons. First, the core determinants of behavior did not remain or stay constant over several decades. Second, commercial computing and Internet markets give rise to a variety of experiences that defy any single characterization.

Nonetheless, the question is worth asking because computing plays such a large role in the economy. Changes in computing now touch both the personal and professional lives of the vast majority of the work force. The basic experience of business computer users has undergone significant change over the last five decades. Starting from a small base of businesses in the 1950s, computing has diffused widely. In 1990, nominal investment in Information Technology (IT) goods totaled $131.5 billion, about 33% of private nonresidential equipment and software investment. By 2000, it was $406 billion and 44% percent.[2]

Similarly, the household experience with computing has also undergone significant change: It began from virtually nothing in the 1970s. Later, a 1995 survey found less than 20% of households had a personal computer (NTIA, 1995). In sharp contrast, an October 2003 survey found that 62% of respondents had a computer at home and an even larger percentage used a computer at work (Mankiw, Forbes, and Rosen, 2005).

These events motivate a wide variety of micro-economic questions about innovative conduct in U.S. commercial computing. In the first section of this review, I discuss these questions in light of six propositions commonly found in the themes of many studies:

- While technical frontiers in computing may stretch due to events reasonably described as "technology-push," a more substantial amount of valuable innovation arises endogenously in response to market incentives and market-oriented events;

[2] See Doms (2004). After falling to lower levels in 2000 and 2001 and 2002, these levels came back up to almost the same levels in 2003 and for the next few years.

- The diffusion and development of computing resembles diffusion and development of a general purpose technology, and as with such a technology, substantial costs arise from creating value by customizing the technology to the unique needs of users;
- The presence of computing platforms shapes incentives to innovate, and the unification or division of technical leadership shapes distribution of value within and between platforms;
- Leading incumbent firms and new entrants face differential incentives to innovate when innovation reinforces or alters market structure;
- Market-based learning activity plays an essential role in innovative conduct, especially in enabling exploration of multiple approaches for translating the frontier into innovative and valuable goods and services;
- The localization of economics activity leads to a concentration of some types of innovative conduct in a small set of locations.

These propositions highlight the continuity between different eras through the underlying economic links between market structure and producer conduct. Indeed, the central thesis of this review highlights continuity, not change. Despite dramatic changes in outcomes, in the predominant product markets, and in the identities of leading sellers, the conditions of market structure shape innovative conduct in firms from one year to the next and, to a large extent, from one decade to the next, in many of the same economic terms.

In the second section, I closely examine the U.S. commercial Internet experience in the 1990s—an analysis that illuminates the strengths and weaknesses of the established frameworks highlighted in the first section. From the outset of the commercial Internet many of its participants have maintained a strong sense about their exceptional nature, as if innovation within the existing value chain for the Internet defied established archetypes of innovation. That view raises a question about whether innovation within the Internet can be assessed with the same economic concepts used elsewhere in computing.

This essay will largely argue that it can be. While the peculiar events that led to the invention of the commercial Internet explain some of the salient and unique features of the commercial experience, much innovative activity resembles conduct seen for many decades in computing markets.

In demonstrating the continuity of economic links between market structure and producer conduct, however, I ultimately achieve almost the opposite—isolating a small set of unique economic factors from the recent era. I account for many of the novel aspects of innovative conduct in Internet markets with three interrelated propositions:

- Innovative conduct related to the commercial Internet did give rise to platforms, but it also gave rise to markets characterized by an extraordinarily high division of technical leadership. In turn, that resulted in an unprecedented dispersion of uncoordinated innovative conduct across a wide range of components affiliated with the Internet;
- Commercial Internet markets involve new organizational forms for coordinating firms with disparate commercial interests, such as open source platforms. Their presence and successful operation accounts for some salient unanticipated innovative conduct;
- The aspirations of entrepreneurs and incumbent firms in commercial Internet markets touched an extraordinarily large breadth of economic activity.

Throughout this review, the narrative will contain a slant towards events in the United States. This slant requires an explanation, since the computing industry today, and especially the commercial Internet, has reached a global scale in operation and in final service markets. This geographic bias partly reflects a pragmatic choice, choosing to compare changes over time, not across geographies. Though comparison between countries (e.g., Japan, UK, or Germany) can be done, such a comparison would widen the scope of the review and take it into too many topics. This review takes an approach that facilitates comparisons over time, concentrating on early computing firms and early commercial Internet participants who substantially, though not wholly, located in U.S. regions. While that focus helps sharpen the contrast between distinct eras, it necessarily limits the scope of the review. It leaves open many comparative questions about the determinants of innovative behavior, i.e., whether the propositions discussed here continue to usefully describe economic events outside of the U.S. boundaries.

II. Innovation in Commercial Computing

From its military and research origins in the late 1940s, computing spread into the commercial realm and has since grown to include an extraordinary range of economic undertakings and a large fraction of U.S. economic activity.[3] Many economists believe this expansion of applications for computing has been a driver of economic growth.[4] Many economic factors have shaped that movement. I begin with explanations that emphasize “technology push.” Finding this approach inadequate in many respects, I inquire about other views that highlight the relationship between market structure and exploratory behavior.

II.1. Stretching the technological frontier and technology push.

In popular discussions, advances in computing have become almost synonymous with advances in microprocessors. This is due to a 1965 observation by Gordon Moore, who co-founded and eventually became chairman at Intel: He foresaw a doubling of circuits per chip every two years. This prediction about the rate of technical advance later became known as “Moore’s Law.” In fact, microprocessors and DRAMS (dynamic random access memories), have been doubling in capability every 18 months over the last three decades.[5]

A similar pattern of improvement—though with variation in the rate—characterizes many other electronic components that go into producing a personal computer (PC), server, or other equipment complementary with computers in many standard uses. This holds for disk drives, display screens, routing equipment, networking and communications equipment, operating systems, communications software, central switches, mainframes and microcomputers, storage devices, input devices, routers, modems, handheld devices, and Internet service provision, to name a few.[6]

---

[3] Many authors have traced the long arc from research origins to commercial form. For overviews, see .e.g., Flamm (1987), or Aspray and Campbell-Kelly (1996), among others.

[4] See the contrasting views of e.g., Gordon (2000), Jorgensen (2001), and Stiroh (2002).

[5] Moore’s law has a long history, beginning with Moore (1965). See Flamm (2003) for a detailed analysis of the underlying components.

[6] See Jorgenson and Wessner (2005) for an extensive review of these many changes.

Indeed, in virtually all applications in electronics, estimates have found extraordinarily rapid rates of improvement in the price per unit of quality of computing, no matter how it is measured. It is also common to measure increases in the ranges of new qualities—that is, to find increases in the number of new qualities provided. That is a robust finding, manifest across a range of computing equipment.[7] For example, in Trajtenberg's (1990) study of computer tomography, the cost of providing a basic scan declined dramatically. In addition, with each year the scanners increased their resolution and their ability to perform new services, achieving milestones that previously had not been possible at any price.

The constant improvement in performance supports the view that many changes in computing arise from "technology push." That is, the invention pushed out the technical or scientific frontier, leading other commercial actors to search for valuable uses. There is a grain of truth to this view, but it also requires proper qualification.

The supporting evidence is well known. Numerous prototypical technologies in computing found their way into products and services long after their invention. These inventions arose in university or commercial laboratories, sometimes as a by-product of basic scientific research goals and sometimes with no direct vision about their application to a valuable commercial activity. Then, these inventions spread through academic papers, by licensing of patents, or the movement of computer scientists and engineers into companies.[8]

Of these inventions, many arose from prototypes built with large subsidies from government funding. For example, the original investment by DARPA (Defense Advanced Research Agency) in the fundamental science of packet switching did not lead to any immediate practical commercial products. Years of sustained funding, however, led to a set of events that broadly subsidized the invention and operation of the basic building blocks of the Internet, such as the experiments that led to the definition of the

---

[7] There is a well-established literature on using hedonic price estimation to measure the rate of improvement in prices. See e.g., Triplett (1989) for a review of these estimates on large systems. See Berndt, Griliches, and Rappaport (1995) for estimates for PCs. For more on similar trends in semi-conductors, upstream to most computing equipment, see Aizcorbe (2006), or Aizcorbe, Flamm and Khursid (2007).

[8] See accounts in, for example, Flamm (1987), Langlois and Mowery (1996), Waldrop (2001), or the overview in National Research Council (2003).

Transmission Control Protocol/Internet Protocol (TCP/IP) stack and its practical implementation in a working communications and computing network. This funding occurred long before the commercial Internet was operational.

Indeed, for the many years the Internet was an engineering novelty, a fascinating invention used primarily by a small group of technically adept networking researchers.[9] Yet, sustained government subsidies helped push out the frontier over time. After the National Science Foundation (NSF) established the NSFNET, government subsidies funded new Internet connections around the country, helped increase its size, and rationalized operations so comparatively unsophisticated users could make use of the network. Only after widespread adoption, did the impact on a substantial amount of research activity begin to become manifest.[10]

While these examples illustrate a role for technology push in conceptualizing innovation in computing, they give the misleading impression that Moore’s Law and related phenomena are exogenous. It is as if breakthrough innovation in computing develops in a deliberate and sequential path, as if each starts as an invention, develops into a prototype, then finally morphs into a valuable product and service. On the surface this conceptualization must be false. After all, much of the behavior underlying the outcomes labeled as Moore’s Law comes from firms with commercial motives, where the managers act in their own interest and in the interest of their stock holders and other providing financing, who expect a return on their investments. These firms push out the frontier (in accordance with Moore’s Law) because it serves their commercial interests, not because they have any desire or strategic interest in supporting industry-wide gains from widespread technological push.

In practice, facets of computing technology and business to co-evolve as researchers and designers push forward the understanding about the costs and

[9] For more on the state of worldwide networking prior to the commercialization of the Internet, see e.g., Quarterman (1989). For more on the origins of the Internet outside the US, see Mowery and Simcoe (2002b).

[10] For example, by some measures, this communications technology had little impact on the conduct of science until the late 1980s. Agarwal and Goldfarb (2006) trace the impact of Bitnet, which was the predecessor to the Internet, on the co-authoring behavior of research engineers. While packet switching had been invented and implemented years earlier, Bitnet began diffusing to universities in the early 1980s, and had little impact on co-authoring behavior until the mid 1980s, and until later, most of the impact was incremental.

commercial value of achieving distinct technical outcomes, as embodied in products and service. Hence, it is more useful to conceive of much innovation taking place in the context of its anticipated effect on the use of computing in on-going economic activity, to analyze how market competition and feedback from user communities shape the direction and rate of commercial innovation. In short, technology push frameworks fall far short of yielding useful insights for economic analysts, managers, and policy makers.

There are numerous alternatives to technology push. We begin by conceptualizing computing as a general purpose technology (GPT).

II.2. Computing as a General Purpose Technology.

Bresnahan and Trajtenberg (1995) define a GPT as a capability whose adaptation to a variety of circumstances raises the marginal returns to inventive activity in each of these circumstances. GPTs are associated with high fixed costs to inventing the technology and low marginal costs to use and re-use. This cost structure both (1) generates heavy early investment—which can occur before and during diffusion of the technology— and (2) leads to frequent repurposing of focal inventions. Rosenberg (1976) describes as "the introduction of a relatively small number of broadly similar production processes to a large number of industries."[11]

The widespread use of computing can be interpreted as evidence of the first characteristic of a GPT—its extensive diffusion. Computing, or IT, ranges from aiding the automated tracking of transactions (a function necessary for automating billing, managing the pricing of inventories of airline seating, and restocking retail outlets in a geographically dispersed organization) to facilitating the coordination of information-intensive tasks, such as dispatching time-sensitive deliveries or emergency services. In addition, the second characteristic of a GPT is evidenced by computing's ability to improve the performance of, for example, advanced mathematical calculations, a function that is useful in activities so diverse as calculating the interest on loans and generating the estimates of underground geologic deposits. As a continued example of the repurposing

[11] The quote is from Ames and Rosenberg, the chapter on machine tools, in Rosenberg (1976). As noted in several essays in Rosenberg's book, these ideas have a long history in the studies of technology, and apply to many more industries than computing.

of focal inventions, note that a closely related function to advanced mathematical computations is computer-aided precision, which improves the efficiency of processes ranging from manufacturing metal shapes to the automation of communication switches.[12]

The creation of intermediate goods like software or networking also shapes the valuation of adopting new IT applications and capital goods. Because it is used in business organizations, IT is deeply embedded in business processes. Accordingly, the business use of IT involves mutual adaptation between business processes and technology, an effort that can lead to large adjustment costs and slow learning about the most efficacious organization for linking inputs and outputs.[13] Part of the complexity arises from the variety of applications for computing in modern economic activity. There may be sharing of noncapital investments across a wide array of processes, and, though the unit costs of sharing are lower for large organizations, the sharing usually does not occur instantaneously or without high coordination costs.[14]

The adoption and implementation of a GPT also can be costly because they lead to changes in other facets of an organization. For example, a GPT may motivate managers to reallocate decision-rights and discretion inside a large organization (Brynjolffson and Hitt, 2000). This is especially so as local business units adapt IT to their local business processes—such as billing, account monitoring, and inventory management—or to the delivery of local services—such as retail sales, the delivery of financial data, and entertainment services. In that case, the boundary of the organization may change along with the adoption of IT, making a direct connection between organizational performance and IT difficult to trace.[15]

---

[12] See Cortada (2003) or McKinsey Global Institute (2001) for an analysis of a wide variety of applications.

[13] For an economic historians' perspective, see David (1990). Analysts from the information systems research community have highlighted determinants of these costs. See, e.g., Attewell (1992), Fichman and Kemerer (1997), Bresnahan and Greenstein (1997), and Forman (2005). Forman and Goldfarb (2006) contain a summary of these factors as they applied to Internet related technologies.

[14] For an illustration in the adoption of machine tools, see e.g., Asterbro (2002, 2004)

[15] This phenomenon results in the value of these investments manifesting in intangible investments. See e.g., Brynolffsson, Hitt, and Shinkyu (2002).

Nevertheless, researchers have tried to trace the links between investment in IT and productivity gains to an organization and the economy as a whole.[16] In general, appropriate microeconomic statistical evidence is difficult to find because such data must contain appropriate statistical variation at a sufficiently small unit of economic production.

In this sense, Atrostic and Nyogen (2005) have a rare and valuable study. They find evidence that use of networking technologies helped raise productivity in a wide array of manufacturing establishments in the late 1990s.[17] Bloom, Sadun, and Van Reenen (2007) also identify an interesting case for isolating the effects of IT use on establishment productivity, finding such evidence among U.S. firms that acquire British firms. These establishments reinvest in their IT operations *after being acquired,* and that leads to measurable productivity advance.[18]

The economics of IT suggests that this evidence should be rare. Aside from simple lack of data with sufficient variance, even with its presence one might not expect a linear connection between changes in input and outputs. Computing frequently enables the invention of entirely new services and products that may or may not provide permanent or temporary competitive advantages.

Moreover, the changes may not be straightforward to measure in terms of value-added or total factor productivity. When new services are reasonably permanent, a private firm may see returns to the investment in the form of increases in final revenue or other strategic advantages. If a new product or service is quickly imitated by all firms, it quickly becomes a standard feature of doing business in a downstream market. The benefits from the new technology are rapidly passed on to consumers in the form of lower prices and better products. In this case, the benefits to a firm do not appear as an increase in revenues; but they exist nonetheless, in the form of losses avoided by the businesses in question.

II.3. Co-invention Costs and Creating Value from GPTs

[16] For a review of the micro- and macroeconomic evidence, see e.g., Jorgenseon, Ho and Stiroh (2005) Or Draca, Sadun, and Van Reenen (2007).
[17] See also Sang and Atrostic (2006).
[18] For related evidence, see also Harrison, Griffith, and Van Reenen (2006).

Co-invention costs are the various costs affiliated with customizing a technology to particular needs in specific locations at a point in time. A competitive supply of tools for co-invention activity can help lower but can never eliminate such costs. These costs shape the ultimate economy-wide cost from deploying and adopting a GPT.

The co-invention costs are frequently difficult to monetize, manifesting, as they do, as lost output, diversion of resources, or disruptions in routines and other "internal costs" inside an organization. Indeed, these should be higher with process improvements that lead to dramatic rearrangements of routine tasks, where firms must self-insure against unanticipated costs. As illustrated by David (1990) in his analysis of the electric dynamo inside manufacturing, computing was far from the first GPT to face such large adjustment costs from disruption of operations.

Bresnahan and Greenstein (1997) hypothesize that co-invention costs are at large scale user installations driven up by complex or idiosyncratic organizational needs, which interfere with the use of generic solutions. They analyze the transition from mainframes to client-server architectures within establishments that already have mainframe computers. Their analysis provides a window on the factors that slow down adjustments inside an organization, because their sample includes many of the heaviest users of computing at the time and many of the establishments that initially adopted computing for business processes. Their findings emphasize the importance of the costs of inventing new uses for computing and adapting it to idiosyncratic and/or complex settings. Such costs slowed down the diffusion of a new technology, often to the users who could generate the highest benefits.

Co-invention costs are not borne solely by users during the deployment of a new GPT. Suppliers may incur them, as part of a strategic approach to developing a new service, limited by the idiosyncratic features of their own organizations and the market niches they serve. As GPTs diffuse, firms explore new ways to make viable businesses from providing services. In the case of dial-up Internet access, for example, the first generation of Internet service providers (ISPs) faced comparatively low incremental co-invention costs because they were complementary to the telephone system and they borrowed many practices from the bulletin board service market. Many firms quickly began generating revenue with just incremental action. Others pursued a variety of complementary businesses, trying to structure business to thrive or merely survive.[19]

An important open question concerns the presence (or absence) of "technological biases" in market outcomes as a result of firms supporting the adoption of innovating computing. In the simplest theory of an *unbiased* technology, its adoption alters the scale of production, without altering the proportion of inputs necessary to achieve that scale. In contrast, a biased technology may manifest itself in several ways, and, as a result, it is not straightforward to observe such biases.

For the first decades of computing, it was generally thought that computing favored substitution away from organizational forms that use less skilled labor, as in the canonical "labor-saving device." This perception was fostered by several popular images, such as the replacement of assembly line workers by robotic tools, or the replacement of a banking teller by an automatic teller machine.

The reality was more subtle than popular images suggest, and the diffusion of the PC made that apparent. More recent research tends to highlight skill-biased technical change.[20] That is, adoption can coincide with a disproportionately more intensive use of highly skilled labor. The returns to higher the value-added associated with more computing disproportionately goes to skilled labor instead of unskilled labor.[21]

Observing the effects of such substitution will be difficult if other aspects of the organization (such as the scale of operations, or valuable features of the end product) change at the same time. It is made further complicated if investors observe these changes and favor some types of investments over others, altering the implicit equity-based opportunity cost of making investment.[22]

Several micro-studies of IT adoption illustrate the complexity of tracing the connection between adoption of IT co-inventions and performance. In the trucking industry, for example, computing altered the allocation of transportation of goods.

[19] Greenstein (2000) explores such costs during the early growth of the Internet. The findings highlight that local conditions shaped the returns to these actions, with firms located in urban areas displaying a great variety of products and services.

[20] For more on such biases, see e.g., Caroli and Van Reenen (2001), Bresnahan, Brynjolfsson, and Hitt (2002), Acemoglu,et al (2007), and Beaudry et al (2006) for PCs in particular.

[21] See.e.g, Beaudry, Doms and Lewis (2006), Autor, Levy and Murnane (2003), Card and DiNardo (2002), Krueger (1993).

[22] A rather provocative discussion of this last point about market valuation, see Shiller (2000)

Hubbard (2000, 2003) examines the use of computing technologies to monitor the performance of trucks. Two distinct applications emerge as valuable, namely, tracking trucks in real time and auditing features of their performance after a completed task. If productivity increases come primarily from the truck being full more often—that is, facilitating matching of truck to prospective tasks, filling backhauls or mixing partial loads—then it is possible to identify the role of computing in bringing about this improvement.

The introduction of on-board computers also improves the ability to coordinate assets in different locations. On-board computing can improve monitoring of driver actions, leading to benefits such as reduced truck depreciation. Yet, fully realizing these improvements requires a rearrangement of the ownership of assets, so tracing all lines of causality in this industry remains challenging.[23]

Another example is how an IT improved the productivity of emergency response services. Athey and Stern (2002) examine the application of computing to emergency services, where value arises from giving timely information to dispatchers of ambulances (in their example). They focus on the consequences for heart attack patients. In particular, the enhanced 911 system (E911) allowed emergency dispatchers to pinpoint the location of a caller. Since timeliness is the key factor in emergency response to probable cardiac events, providing accurate information enabled a more rapid response. In this way, the introduction of E911 was shown to reduce the probability of a patient arriving with a high-risk pulse rate as well as the probability of mortality within 48 hours.

Research tends to focus on the factors that drive up co-invention costs, highlighting the myriad reasons why new technologies do not become employed rapidly. Yet, co-invention costs do not have to be high, and this can form an important part of the explanation for rapid diffusion of a GPT. For example, Forman (2005) examines the early adoption of Internet technologies at 20,000 commercial establishments from a few select industries. He concentrates on a few industries with a history of adopting frontier Internet technology and studies the micro-economic processes shaping adoption. He finds widespread use and adoption of basic technologies, such as email and browsing, consistent with low co-invention costs for networks supporting these applications.

[23] For example, see the analysis of Baker and Hubbard (2003, 2004).

Later, Forman, Goldfarb and Greenstein (2003a) extend Forman's work to all non-farm private establishments in the U.S. economy, surveying establishments with over 100 employees, where accurate and extensive data about the use of the Internet exist for roughly half the establishments of that size. Projecting from this survey to the economy as a whole, they estimate that by the end of 2000 close to 90% of all such establishments had access to email and browsing. Some industries had reached saturation, such a printing, parts supply, and many financial activities; other industries had high rates of adoption (over 80%), while others such as waste management, garden supply, and social assistance had lower rates. They conclude that co-invention costs were low for almost all industries because (1) PCs had already diffused, (2) supply of related services was widely available across the country, including low-density areas, and (3) the incremental costs of adopting these additional activities involved only a small set of steps.

Forman, Goldfarb, Greenstein (2005) examines the influence of factors on the marginal adopter. Holding all else constant, they show that business participation in the Internet is more likely in rural areas than in urban areas. This is particularly true for technologies that involve communication across establishments. Nevertheless, talk of the dissolution of cities is premature. Frontier Internet technologies for communication within an establishment appear more often at establishments in urban areas, even with industry controls. The difference between marginal and average rates is largely explained by differences in industry composition across major cities. More IT-intensive industries tend to cluster in urban areas. The effects of urban leadership and industry composition interact in a complementary way for advanced applications and that interaction exacerbates agglomeration in use.

While co-invention costs were low for basic browsing, they were much higher for any significant investment in enterprise computing. For example, building ERP systems consistent with Internet protocols and integrating them into enterprise operations involved extensive customization to reflect each firm's production process, reporting norms, security and accounting procedures, and supplier relationships. Such investment could either provide managers with additional discretion or serve to centralize authority within firms.[24] Similar issues shaped a wide range of enterprise level investments in Internet-enabled applications for procurement, distribution, inventory tracking, coordination of payroll, and so on.[25]

While co-invention costs help classify the costs of learning how to turn a GPT into something useful, analysis of innovative conduct in computing often requires something else. It needs to analyze the effects of specific institutional details on innovative conduct, such as the goals of the users, the procedural patterns of the industry-wide organizations, and the identity of the leading firms. We begin to develop that analysis next.

II.4. Platform Competition Shapes Economic Incentives

The direction of innovative opportunities is shaped by the relationships between firms, and those relationships are shaped by the presence of platforms. In any given era, computing markets are organized around platforms—a cluster of technically standardized components that buyers use together to perform the aforementioned wide range of applications. Platforms shape the incentives to pursue directions of innovative activity and, arguably, also its rate.[26]

Such platforms involve long-lived assets, namely, both components sold in markets (i.e., hardware and some software) and investments made by buyers (i.e., training and most software).[27] Important computing platforms historically include the UNIVAC, the IBM 360 and its descendents, the Wang minicomputers, IBM AS/400, DEC VAX, Sun SPARC, Intel/Windows PC, Linux, and, recently, TCP/IP-based client-server platforms linked together.

---

[24] See, e.g., Brynjolfsson and Hitt (2000), or Bloom and Van Reenen (2007).

[25] For further review of these and related studies, see Forman and Goldfarb (2006).

[26] The word "platform," as used in engineering, is a different notion. In this context, the emphasis is on the formation of valuable interconnected economic relationships and how their presence alters incentives to undertake new innovative activity. See e.g., Bresnahan and Greenstein, 1999, or Cusumano and Gawer, 2002, among others.

[27] Within the computer industry, the user investments are often branded as "sweat equity" by sales forces who must be cognizant of their presence to make a sale. See the contrasting descriptions of Cordata (2003), and Shapiro and Varian (1997), where the former uses the industry vernacular and the latter uses the economic theory of "switching costs" to discuss much the same phenomenon.

Vendors tend to sell groups of compatible products under umbrella strategies aimed at the users of particular platforms. In the earliest eras, the leading firms integrated all facets of computing and offered a supply of goods and services from a proprietary source. In later eras, the largest and most popular platforms historically included many different computing, communications, and peripheral equipment firms, software tool developers, application software writers, consultants, system integrators, distributors, user groups, news publications, and service providers. While some of these might take actions to serve proprietary interests, they all commit to the platform, and invest with the expectation that the platform will continue.

Platforms display a form of increasing returns that is sometimes given the labels "network effect" or "bandwagon effect."[28] That is, the value of participating in the platform grows as more participants commit to it. These benefits accrue through a variety of mechanisms: Users may benefit from participating in a large platform because large platforms display larger selection, lower prices, and greater opportunities to "mix and match" components from multiple suppliers. Vendors may benefit from participating in larger platforms because it provides them access to thicker demand for niche products and more accurate perceptions about the long-term viability of accumulated groups. Larger platforms also allow firms to specialize in innovating on a few areas while leaving other markets to specialists in other complements.

The emergence of platforms tends to coincide with the emergence of a standard bundle of components. That is, a standard bundle embodies a set of common arrangements of components for delivering services. Most users of a platform are similar in this respect.

This explanation also begins to hint at why some users and participants avoid platforms. The standard bundle may constrain the functionality. So, for example, technically adept users tend to favor different standard bundles than users that make up a mass market. Similarly, a standard bundle constrains a vendor's ability to differentiate.

[28] Jeffrey Rolhfs is credited with the earliest models of this phenomenon, which he motivated on his observations about the failed videophone at AT&T. Rolhfs (2001) contains numerous case studies, as well as an intellectual history of his thinking about bandwagon effects. Katz and Shapiro (1985, 1986) present a model of networks with endogenous pricing, calling this a "network externality." Their models consider a variety of settings where those externalities are internalized or not. Also see Farrell and Saloner (1985) for a related definition of network effects. For an overview of this approach, see Katz and Shapiro (2000).

For strategic reasons, occasionally vendors will try to break with standard bundles to achieve such differentiation.[29]

Until the early 1990s, platforms helped define the margins between most market segments, which were distinguished by the set of common functions in which a group of sellers/users shared an interest. These segments represented clusters of technical skills at firms and clusters of operations at users. Typically these shared interests corresponded to the size of tasks to be undertaken and the technical sophistication of the typical user. Mainframes, minicomputers, workstations, and PCs in decreasing order, constituted different size-based market segments.[30]

The most popular platform in the late 1980s and 1990s differed from the prominent platforms of earlier years. For example, the workstation appealed to technically sophisticated users, and typically employed advanced microprocessor power and some modified variant of a Unix operating system. Numerous companies competed with proprietary versions of hardware and software designs. Eventually the leading firm in this segment became SUN Microsystems, which employed a mix of proprietary and nonproprietary technologies that appealed to users and a large community of software application developers.[31]

The other popular small system was the PC. It began in the mid-1970s as an object of curiosity among technically skilled hobbyists. After a brief period of competition among designs within the segment, it became a common office tool after the entry of IBM's design. Unlike prior computing platforms, this one eventually has diffused into both home and business use. From the beginning, this platform involved thousands of large and small software developers, third-party peripheral equipment and card developers, and a few major players.[32]

---

[29] Bresnahan and Yin (2007) provide an insightful analysis of the strategic imperative to achieve such differentiation in a standards battle and the factors that shape success and failure.

[30] See Bresnahan and Greenstein (1999) for a reinterpretation of computing history continuity and change in terms of its platforms.

[31] See e.g., the account by Baldwin and Clark (1997), and an update in Baldwin and Clark (2006).

[32] These events are well known and well documented. For example, see the accounts in Cringley (1992) or Freiberger and Swaine (1984), among many. For a statistical study of this competition, see Gandal, Greenstein, and Salant (1999). Bresnahan and Greenstein (1999) analyze these events in terms of platform economics.

More recently, control over standards has completely passed from IBM to Microsoft and Intel. Microsoft produces the Windows operating system and Intel produces the most commonly used microprocessor. For this reason the platform is often called *Wintel*.

The networking and Internet revolution in the late 1990s is responsible for blurring prior familiar distinctions. At first, these new technologies involved a combination of work-stations and PCs hooked together with a local area network (LAN). These innovations made it feasible to build client-server systems within large enterprises and across ownership boundaries. Firms with dominant positions in the earliest client-server platforms included Novell, 3Com, Oracle, and Cisco.

Before client-server systems completely diffused to all enterprises, another innovation altered the path of development, the Internet. As a technical matter, the Internet involved a series of standard protocols that permitted a user to move data within and across networks as long as those networks employed the same protocol.

There were many new features to the commercial Internet, but two features especially stood out as a type of commercial computing network technology. First, the Internet was designed to have its intelligence at the end of the network. That is, users had to adopt applications in the PCs and workstations that were compatible with one another, but did not have to worry about any of the devices or protocols inside the network.[33]

Second, once the commercial Internet had diffused (by 1997 to all major cities in the United States), a remarkable set of new possibilities emerged: The Internet made it possible for users and vendors to move data across vast geographic distances without much cost, either in operational costs and/or in advanced set-up costs of making arrangements for transport of data. Together, those two features enabled enormous combinations of users and suppliers of data that previously would have required bilateral—and, therefore, prohibitively costly—agreements to arrange. In brief, it enabled a network effect where none had previously existed, involving participants who could not have previously considered it viable to participate in such a network.

[33] For an account of how this built up in email design, see Partridge (2008). For an accessible discussion of the design principles behind the Internet and why recent events threaten their continuance, see Blumenthal and Clark (2001).

Today such networking employs Internet-based computing systems connected across potentially vast geographic distances. This results in the emergence of a "network of networks,"[34] which employs a mix of nonproprietary designs, from such organizations as the IEEE (Institute of Electrical and Electronics Engineers), the IETF (Internet Engineering Task Force), the World Wide Web Consortium (W3C), and others that will be describe in more detail in later sections. And these organizations co-exist with providers of proprietary products and services. The latter are sponsored by firms such as Microsoft, SAP, Oracle, Google, Yahoo, who compete with each other, and use software built around their own designs, as well as around those from open sources, such as Linux, Apache, and MySQL, as well as others described in more detail in later sections.

II.5 Innovation Within and Between Platforms

Platforms have existed in every era of computing. Platforms are significant because the direction of innovative activity takes place within the constraints of platform competition. Competition between platforms arises whether or not firms with proprietary interest (in a platform) either wholly or partly control the platform. Innovative activity contributes to altering the conditions of competition either between or within platforms. There have been few empirical studies of historical competition between platforms, primarily because such competition is infrequent. New proposals for platforms rarely develop past conceptualization into a commercial form that users will buy. In addition, once they are widely adopted, existing platforms tend to be hard to stop or slow down. As a result, markets do not give rise to many settings where several viable platforms last for long and compete side by side.

Many of the studies of the mainframe era, for example, attempt to understand the rise of the IBM System 360 as the dominant computing platform. The interest is understandable, since this was the most lucrative commercial innovation in computing for several decades. The CEO (chief executive officer) at that time, Thomas Watson Jr., led a huge development effort, putting at risk virtually all the assets of the firm. This family of

[34] See, e.g., Noam (2001) for studies of "networks of networks." For a range of writing about the operations of this network and the numerous issues raised in trying to manage this at a large scale, see e.g., McKnight and Bailey (1998), Kahin and Keller (1995, 1997), Mansell and Steinmueller (2000), Compaine and Greenstein (2001), Cranor and Greenstein (2002) and Cranor and Wildman (2003).

products succeeded in becoming the primary platform for the development of numerous office-computing applications in banking, payroll, inventory accounting, and many other key innovative applications in the 1960s and 1970s.

Several studies have sought to understand either the circumstances the led to IBM's development or the consequences of its success. Katz and Phillips (1982), for example, focus on what led to IBM's success by highlighting the learning needed to change organizations and the product designs to accommodate the changing needs of customers. They find that large-scale computing had begun as a scientific and military pursuit, with most development funded by governments. The early technical leaders of computing for scientific and military applications had difficulty anticipating the operational requirements and redesigns valued by commercial users with different needs. The emergence of mass-market office users for computing changed the value of learning about commercial users, thereby altering the value of the firms that had been early pioneers.

Fisher, McGowan and Greenwood (1983) focus on the consequences of IBM's success, by advancing a view that emphasizes the dynamic changes in the market place and the rewards to entrepreneurial commercial initiatives at leading firms, in this case, primarily IBM. They credit IBM management with creating and operating processes that translate the information gained from learning about new technologies and user experience into new product designs.[35]

There are two approaches to understanding how competition works within platforms. Consider first how the links inside a platform shape the conditions of innovation. Buyers and sellers become linked by their technical interdependence and their continuous economic relationships. As a result, it is common for these ecosystems of suppliers and users to evolve and revolve around a common bundle of components. Innovative opportunities and constraints thus depend on what new offering appears that departs from that standard bundle.

The ecosystems of vendors and users differ in their stability, turnover, and size, depending on the requirements of the services performed.[36] For example, consider the

[35] See also Fisher, McKie and Mancke (1983).
[36] See the extensive analysis in Messerschmitt and Szyperski (2003).

difference between the ecosystem for providing security software and that for providing printers, each of which are common components in a standard PC purchase today. The security market involves firms with a labor force that constantly works to predict the behavior of hackers and is potentially on-call and ready for immediate crisis when a virus spreads. In contrast, the printer markets contains participants more reminiscent of a group of typical manufacturing organizations, involving regular supply chains for parts and components, third-party distribution, and servicing of end products.

A second approach for understanding competition within platforms tries to incorporate recent circumstances into its method. In general, no firm controls all aspects of a single platform. Competitive analysts focus on understanding how ensembles of participants (within a platform) behave, either in competition with one another or in cooperation.[37] There also are several studies of format wars or standards wars, which may or may not involve large groups of firms providing a standard bundle, as in a platform.[38] The more recent literature also investigates why the era of single ownership over an entire platform ended. This brings us to a discussion of divided technical leadership.

II.6. Divided Technical Leadership and Innovation Conduct

In the 1950s, computing was a novel technology and only a handful of experts understood all its key features. During the past five decades, the dispersion of expertise has transformed dramatically to encompass a vast ensemble of participants from a variety of technical and commercial backgrounds with varying kinds of motives. That transformation coincides with a change in commercial conditions for computing. No longer does a small set of expert engineers understand *all* dimensions of computing. This does not imply that expertise has no value; rather, a small number of experts do not largely determine the rate and direction of technical change in all aspects of computing.

[37] For example, Bresnahan and Greenstein (1997) analyze such competition in the context of the flight between mainframe and client-server for large establishments. Bresnahan (1999) considers some of the barriers faced by "smart server–dumb client" platforms during an early period, offering a model of sequential entry from "one adjacent component market into another." As another example, Gandal, Greenstein, and Salant (1999) examine the competition between different platforms in the early PC market, highlighting the role of feedbacks between components markets.

[38] On the VHS/Beta war, see Ohashi (2003), on the hardware/software interplay in DVD adoption, see Gandal, Kende and Rob (2000), on the DiVX war, see Dranove and Gandal (2003), and on the 56K modem war, see Augereau, Greenstein and Rysman (2006).

In other words, there has been a secular trend toward an increase in the number of firms that possess the necessary technical knowledge and commercial capabilities to bring to market some component or service of value to computing users. If a firm does not possess such capability, it can be easily acquired through market means, such as hiring a small team of qualified engineers. Brensahan and Greenstein (1999) call this feature of market structure *divided technical leadership.*

To illustrate, consider some of the key innovations of distinct eras: In the 1960s, when technical expertise was not widely dispersed, the most lucrative innovation of commercial computing, the IBM System 360, arose from using designers and developers employed entirely by a single firm. It involved redesigning every major aspect of commercial computing. Although many of the inventions had arisen elsewhere, many also came from within IBM. More important, the key invention had not arisen anywhere—instead it involved putting all the other inventions together. Coordinating that large and diverse team pushed IBM to the boundaries of what any ambitious firm feasibly could manage. Its innovations thereafter employed various pieces of the System 360 for new applications, such as airline reservations systems or new account-tracking systems for financial users. Though it faced many imitators, IBM employed unique assets and products in producing, delivering, and servicing its own products, and that helped protect its innovation from imitation.

Consider, in contrast, divided technical leadership, which supported commercial initiatives by those specializing in supply of innovative components. Thus, broadly speaking, at the equipment layer one set of firms specialized in supply, while another distinct set of firms carried the data using frontier operations. Another set of firms brought new storage devices to market, yet another provided frontier software applications for users, and an even different set performed frontier services. This is sometimes called specialization in *horizontal* or *component* layers to distinguish it from competition between integrated systems.[39]

The specialization of supply frames one of the distinctive strategic issues of the modern era. Firms with quite different capabilities, specializing in one or a small set of

[39] This was famously summarized by Andy Grove, CEO of Intel, who described the distinct layers of the PC industry. See Grove (1996).

components, cooperate with others at the boundary of their respective firms. In personal computing, for example, an array of distinct firms arose that specialized in supplying different parts of the PC (e.g., many firms provided the electronic components), while different firms provided the software. An entirely different set distributed the final product and became involved in servicing it. The benefits of allowing users to mix and match components and service outweighed most of the benefits of coordinating production entirely inside one firm.

Markets with dispersed technical leadership tend to contain a variety of firms. Here *variety* means the following: Firms use different commercial assets in different locations, different personnel with distinct sets of skills, different financial support structures with different milestones for measuring progress, and even different conceptual beliefs about the technical possibilities. In this sense, variety shapes the conditions of competition. One firm's assessment of the returns from innovating does not need to be the same as another's. Different assessments result in different methods for achieving the same commercial goals, which may lead to different costs, or different commercial goals altogether, such as targeting different customers.

Divided technical leadership also reduces barriers to entry in component markets, which, in turn, supports widespread availability of potential suppliers of a component or service. Thus, entry of new component firms into a market arises from one of three paths:

(1) A component specialist in a one technical "area" may develop expertise on a "neighboring" area in the course of operations – for example, a hosting firm will learn plenty about security software. That firm may use existing personnel and assets in that area to develop another specialty component that extends its existing business;

(2) An entirely new entrant may arise, starting from scratch, assembling components made by others or hiring technical talent for a newly focused goal or newly re-optimized organizational form. Such firms often get financial backing from venture capitalists and view their goals in terms of a race to achieve a functional leap over all other firms in a specialized area;

(3) A firm with a broad proprietary platform, such as Intel or Microsoft, may seek to enter by embedding similar functionality in one of its existing products, either through acquiring another specialty firm (who entered as 2) or developing their own proprietary version (as in 1).

Such competitive conditions can lead to a reduction in any specific supplier's ability to possess a unique set of assets or employ all the innovative experts in the world. Where once firms used technical prowess for commercial gain, now firms find their technical skills *commoditized*, whereby the firm becomes a supplier of a product (technical prowess) that is available from many vendors. These issues shaped the most lucrative computing innovations of the mid-1990s.

First consider Microsoft's Windows 95—an innovation that involved redesigning every major aspect of the operating system for PCs and pushed Microsoft to the boundaries of what its organization feasibly could manage.[40] Microsoft embedded many functions inside Windows 95 that had previously been provided by specialist software firms. Despite its internal capabilities, Microsoft's managers deliberately left an extraordinarily large amount of commercial computing untouched, especially in equipment and application markets where Microsoft had no products or strategic interests. This was a rational decision that recognized both the limits of its own firm and the capabilities of others.[41]

Next, consider another lucrative innovation of the mid-1990s, upgrades to the microprocessor at Intel. Intel's managers repeatedly faced decisions about whether to initiate new projects or reinvest in existing projects, some of them in areas "on the motherboard of the PC," where this functionality complemented the microprocessor. When they reasoned that such investments would expand final demand for PCs, then they would invest, as they did, for example, in redesigning the bus for the PC.[42] Throughout the 1990s, managers chose to invest in interfaces that worked directly with their existing

[40] See e.g., Cusumano and Selby (1995).
[41] These types of decisions are discussed at length in Bresnahan, Greenstein and Henderson (2006).
[42] In their own words, the managers would invest in trying to improve some aspects of the PC if it would "grow the pie for everyone." See the discussion in Cusumano and Gawer (2002).

product line, while avoiding projects where plenty of other suppliers had the ability to reach the frontier at lower cost.[43]

One might summarize it thusly: While competition between platforms determines prices for customers deciding between platforms, divided technical leadership shapes the competition for, and division of, returns within a platform. These two margins differ, and rather distinct aspects of firm conduct determine each of them.

Divided technical leadership, along with the growth of the market to support specialization, also partly explains several other prominent features of the computing industry on a global scale, such as the increasing geographic concentration of some activities in some local areas. As many firms specialize in different components, supply chains for most computing hardware and software products no longer falls under control of one firm. The location of production becomes subject to competitive forces, and the identity of leading firms can change. This trend has been widely seen across many computing components, such as displays, mobile devices, and other hardware devices.[44] In addition, as in other manufacturing processes, the increasing use of sophisticated IT helps coordinate design and production involving firms from many countries and continents, which also contributes to spreading it to many locations. Divided technical leadership also affects software, a labor-intensive activity where coordination and monitoring costs have declined enough to support geographically dispersed production.[45]

Despite the emergence of divided *technical* leadership, computing markets do not display frequent turnover in the identity of *market* leaders. Why is that? I next examine persistence and racing as a means to answer this question.

II.7. Racing and Persistence by Incumbent and Entrepreneurial Entrant

Pricing at some level above unit cost requires something rare among specialized providers—for example, valued brand, frontier features, unique service, or better distribution. Firms take a variety of approaches to developing these assets in technology markets, even if they remain unique for only a short period. Said succinctly, firms face

---

[43] A review of these issues can be found in Gawer and Henderson (2007).
[44] For example, see Dedrick and Kraemer (2005) for a study of the sourcing for personal computers. See Hoetker (2006) on flat panel displays.
[45] See, e.g., Mowery (1996), or Arora and Gambardella (2005).

incentives to fund the search for and creation of innovations that use proprietary assets, such as existing effective distribution channels.

Among these many actions, firms compete with one another to reach a unique technical accomplishment. This type of competition is often labeled *racing*. Racing arises in the midst of upgrades to existing products or during times when firms offer new product introductions. It has been a widely documented behavior, found in disk drives, printers, software, and other components, with considerable variance in the later performance of early winners of races.[46]

Despite frequent and sometimes dramatic technical improvements in the frontiers of technology, many features of the most common platforms in use tend to persist or change very slowly; this is labeled *persistence*. Persistence happens even when the frontier has far outstripped the most common technologies in use.

That observation raises the question: Why do incumbents or entrants *not* take advantage of every technological opportunity? What are the costs and benefits of racing and persistence? As it turns out, sometimes firms either cannot or will not do the same as those who race to the frontier. This is because many durable components make up platforms. Though old technology loses its status as a frontier application as it becomes obsolete in comparison to newer frontier products, it does not as quickly lose its ability to provide a flow of valuable services to users. In brief, a service may be valuable even when it is not on the frontier if it enhances and preserves the value of previous investments while simultaneously giving users access to some new functionality.

A "backward-compatible" upgrade or improvement is one that works with or remains compatible with existing equipment. The label is revealing. A "backward" technology is deliberately not one that moves "forward" to the frontier, but, instead,

[46] For example, Khanna's (1995) study of new product introductions in frontier computing documents such behavior. Stavins (1995) examines the incentives of PC firms to enter the frontier or the "middle" of the product space. Lerner (1997) examines whether hard-disk firms gain premiums in their pricing from developing new frontiers. Greenstein and Wade (1998) examine whether competitive setting shapes the incentives of firms to bring out new products. Cockburn and McGarvie (2006) study how intellectual property shape the incentives of new software entrants; DeFigeiredo and Kyle (2006) look at existing incumbents, principally Hewlett Packard, or at new entrants developing new frontiers of the product space for laser printers. Prusa and Schmidt (1994) present evidence that few of the early entrants into PC software survived thrusts into other segments.

remains compatible with prior functionality. It also creates a demand for support and innovative service activities to reduce the costs of making the transition from old to new.

There is often a fundamental tension between aspirations toward the frontier and a backward-compatible upgrade. While this trade-off always exists, platform leaders face high opportunity costs for mismanaging the trade-off. There may be distinct adequate solutions to a customer's demand for backward-compatible components and for the frontier. Firms also might need different designs in order to satisfy both frontier and backward-compatible customers. The open question is whether most customers would be satisfied with the same design offered by a platform leader.

During the era of the large systems, for example, IBM faced a set of related issues after the success of the System 360. Instead of designing its next mainframe from scratch, it chose to introduce a family of systems, the System 370, which remained backward compatible with many of the investments made by users of the System 360. That design choice led to a lucrative business opportunity with many existing mainframe customers; however, it made it quite difficult for IBM to satisfy certain sub-segments of users, particularly those focused on innovative high-speed computing.[47] Indeed, IBM faced essentially the same trade-off for many years. In each case, IBM always made at least one option available that allowed its present customers, the vast majority of mainframe users, to upgrade without losing prior investments.

IBM's managers did try to extend the reach of its systems to appeal to new users, and they succeeded occasionally in these races and occasionally not. In these cases, IBM's competitors attempted to design systems that appealed to niche users, but the systems were not compatible with the platform. Notably, in races on the low end, where such upgrades were rarer and the demand for backward compatibility weaker, IBM experienced a variety of new entrants and competitors. IBM also tried a variety of different partially compatible systems to compete. None of them did especially well until the PC in 1981 and the AS400 some years later.[48]

A similar dilemma arose in the era of widely divided technical leadership in PCs. After the rise of 386 computing in PCs, IBM's leadership in PCs had diminished. Intel's

---

[47] Accordingly, other platform providers, such as Control Data and Cray focused on them. See, e.g., Fisher et al (1983).

[48] See the analysis in Bresnahan and Greenstein (1999), or Bresnahan, Greenstein and Henderson (2006).

management considered new designs for the micro-processor, particularly whether it should follow the lead of several other frontier firms and break with prior designs in order to reach the frontier. Although such a break would help Intel compete with others that demonstrated superior performance, it would sacrifice some backward compatible functionality. Instead, after considerable debate, Intel chose to follow a path of backward-compatible improvement throughout the 1990s. In retrospect, this was a good decision for the firm's stockholders, even though it was far from obvious as the time.[49]

Aspects of this same dilemma arose during Microsoft's strategic thinking in the mid-1990s. For example, though Windows 95 was a new operating system for the PC, Microsoft eased the portability of software applications that had run on Windows 3.0 and 3.1, which sat on top of another operating system, DOS. Launched in August 1995, Windows 95 included a sub-window for porting files from the DOS environment, where vast majority of non-frontier programs were still surviving. Microsoft went to such lengths for numerous reasons, but among them was the concern that—in the absence of such features—these existing users of DOS would not migrate to Windows 95 but would, instead, form a subgroup of users that another competitor in an operating system market might support and use as a base from which to grow.[50]

These last two examples allude to a subtle, but important feature of why the tension between racing and persistence is so central to the analysis of innovation and platform leadership: The incentives to meet demand for backward compatibility fall asymmetrically on existing and entrepreneurial entrants.

In one common scenario, an entrant seeks to imitate the incumbent with a backward compatible offering that undercuts margins on one component. For example, RCA sought to imitate IBM's System 360 with a backward compatible design, only to find that this was an expensive and ultimately unprofitable activity. As another example,

[49] Grove (1996) describes the vigorous debate around this crucial decision, his initial desire to adopt a frontier design, RISC (reduced instruction set) instead of CISC (complex instruction set). A retrospective view is found in Tedlow (2006).

[50] For example, at that point in time Microsoft had recently experienced competition from OS2, an IBM operating system that closely resembled Windows 30 and exceeded its functionality, and from DR-DOS, an imitation of the DOS operating system.

DR-DOS sought to imitate Microsoft's DOS, only to find that getting access to distribution channels was expensive and difficult.[51]

The innovative outcome from this common scenario depends on the decisions made by incumbent firms. For many years, for example, IBM sought to compete with all entrants, experiencing considerable success in mainframe markets for corporate computing throughout the 1960s, '70s and, '80s. In contrast, it had a more mixed experience in other segments when bringing innovations to market. For example, in the 1970s and 1980s, Digital Equipment Corporation had considerably more commercial success in factory floor computing and general purpose minicomputers. Similarly, in the 1970s, Wang had more success in word processing—until the rise of the IBM PC in the mid-1980s.

The tolerance of early users for less-than-perfect products plays a role in the other common scenario. An entrant may seek to satisfy a segment of users that the incumbent firm's backward compatible offering neglects after establishing its user base.[52] For example, many work station firms in the mid to late 1980s sought to satisfy the needs of users with demand for high-speed computation, a function that the common PC could not meet as efficiently and the mainframe could not provide as cheaply. It turned out that in the early 1990s, these same workstation firms were often the biggest supporters of client-server architectures, when these firms again tried to expand their functionality to fill a need that neither incumbent platform had met satisfactorily (though not for lack of trying).

In short, platform leaders have incentives to expand the scope of platforms from which they profit, and they have incentives to aspire to continuity in the use of that platform. Entrants, in contrast, have incentives to consider whether to commit to an existing platform, or join another that might compete with it. In turn, that translates into high incentives for incumbents to support design of new proprietary standards for an existing platform, but not non-proprietary standards that might lead to more competition between platforms. On the other hand, entrants of applications prefer to make them compatible with as many platforms as possible, which lead to incentives to work towards

[51] For an analysis of the issues Microsoft faced and their behavior, see Gilbert (1998).
[52] For analysis of some of these issues, see e.g. Bresnahan and Greenstein (1997), Bresnahan and Yin (2006), Arora, Caulkins and Telang (2006).

non-proprietary standards, or other technological tools to reduce the costs of supporting cross-platform applications.

The flow of events during more recent experience has also depended on the choice made by incumbent firms. For example, in the early 1990s, Microsoft devoted enormous organizational effort and energy to producing Windows 95, and it reaped enormous profits. At the same time, Microsoft's management misinterpreted events in the Internet, not recognizing how a series of innovations from different corners would lead to the viability of many businesses founded on browser-based computing. Accordingly, the company spent the better part of the mid- and late 1990s and beyond trying to make up for a comparatively late start, while also selling in large quantities Windows 95. Microsoft altered the direction of its investment, supporting some Internet technologies, defended the value of the investments it had already made.[53] Then, after the potential for the alternative platform built around Internet and Web standards diminished, it cut back on many of those same investments.[54]

In each platform, it is rare to observe more than a small number of firms acquiring leadership positions. It is unsurprising, then, that questions about how incumbent firms react to new entry and defend existing positions in valuable markets have attracted antitrust scrutiny. For example, IBM's behavior in peripheral markets attracted such attention in the 1950s and late 1960s. Intel and Microsoft's behavior in the mid-1990s also attracted such attention.[55] Did such attention alter innovative behavior? Almost certainly it did, though precisely how remains open as of this writing.[56]

II.8. Economic Experiments and Market-Based Learning

So far, I have highlighted that innovation in computing markets often does not begin with events in a laboratory nor does it follow a predictable sequential set of stages. Instead, activities outside of a laboratory often take primacy, such as innovations that

---

[53] This story is well documents in many places. See Cusumano and Yoffie (1998) for an account of Netscape's Founding.

[54] See Bank (2001) for primary account of this behavior. Bresnahan, Greenstein and Henderson (2006) provide a summary of Microsoft's actions.

[55] For an overview of differing viewpoints, see e.g., Schmalensee (2000), Fisher (2000), Henderson (2000), or Bresnahan (2004).

[56] See, e.g., Gilbert (2006) and Baker (2007).

arise from market experience. A series of studies of innovation in computing highlight how firms learn from their experience in markets, particularly when market experience alters knowledge about the value of a good or service or the costs for bringing a service to market. Following Rosenberg (1994) and Stern (2004), we label these events *economic experiments.*

Economic experiments involve more than just changing knowledge pertaining to technical invention; economic experiments may also change knowledge about business operations and organization that translate technology into economic value. By this broad definition, economic experiments encompass a wide range of market-based learning. The most common economic experiment is incremental in its technical scope and ambition. It aims at learning lessons with immediate consequences for a business. Though incremental, it can involve decisions of the utmost importance to the business, such as learning information about the pricing for a new service using a new technology.

From one perspective, this activity is mundane and almost routine. Managers would authorize the expenditure of resources, redirect personnel, alter a feature of an existing service, develop a new service, advertise a service or not, and then wait to find out whether these investments paid off in terms of additional revenue, market share, or pricing authority. Failure was not regarded automatically as a waste of resources if it led to valuable learning (for example, a failed small-scale experiment could help managers avoid costly mistakes on a larger scale).

What do market participants learn from their experiments? They gain information that reduces uncertainty about the source of value in markets. From where does such uncertainty originate? Rosenberg's (1995) analysis of uncertainty in computing markets provides a suitable framework for understanding how experience shapes market based learning in computing. In this structure, five factors prevent market participants from forecasting the future:

- primitive technology;
- unexpected complements;
- narrow search which yields applications with unexpected breadth;
- unanticipated systems;
- unpredictable user valuation.

Uncertainty arises because technologies are primitive at the time of introduction. Market participants cannot anticipate the technology's use until it becomes more refined. For example, in 1975, it was difficult to imagine the use for an advanced PC in 1995, because the 1975 processors were incapable of doing more than add numbers. Only experience with faster processors contributed to understanding what the PC could achieve.

Market participants also cannot learn about key complements without experience. One invention motivates searches for complements, a search whose outcome may be difficult to predict. For example, improvements in microprocessors motivated inventions of complementary parts in the motherboard, software, printer devices, screens, and myriad input/output devices, which further motivated improvements in microprocessors. The result from such searches can be learned only after firms introduce new products.

A very subtle barrier to forecasting arises because a narrow search may generate a wider set of applications than were the original motivations for the search. For example, a short-range wireless networking technology—now popularly known as Wi-Fi—did not arise from a single firm's innovative experiment. Instead, there were many potential business applications for this standard; One of the earliest prototypes had been in wireless terminals [57] and another had been in a large-scale LAN for a university campus.[58]

After Wi-Fi was designed in the IEEE 802.11 committee, numerous businesses began directed experiments supporting what became known as *hot spots,* which was an innovative idea altogether for retail provision of wireless computing. A hot-spot in a public space could be free—installed by a home-owner, maintained by a building association for all building residences, or supported by a café, restaurant, or library trying to support its local user base. Or, it could be subscription-based, with users signing contracts with providers, supported (usually with commercial arrangements) by a café', airport, restaurant, or other commercial entity. In any events, a hot spot was a use far outside the original motivation for the standard.

---

[57] Vic Hayes, one of the earliest developers of wireless technologies and standards, and chair of the IEEE 802.11 committee during the 1990s, first developed wireless technologies for National Cash Register (NCR) (a sub-division of AT&T then, today a division of Agere Systems). In that capacity he first developed wireless terminals for stockbrokers. See Krariff (2003).

[58] See the description in Hills (2005) of the beginning development of the equivalent of a wifi network for the Carnegie Mellon campus in Pittsburgh, which starting in 1993.

The Internet is an excellent example another type of uncertainty, namely, unanticipated systems. Market experience is required to recognize the value of a system of complements. Several components may be comprised by a system that delivers a functionality whose value exceeds the value any component could deliver alone. In computing, for example, it was difficult to forecast how the PC, fiber optics, and appropriate software would provide functionality as the Internet. For many users it could not be appreciated until experienced.

Finally, user valuation of end products is difficult to predict and often cannot be understood until experienced. The history of the market is full of user valuations that a mere survey would not have revealed. For example, there were well-known examples of underestimates of mass-market enthusiasm for the back-office mainframe computer, general purpose minicomputer, PC, browsing, mobile computing, and other facets of electronic commerce.

Economic experiments continue to shape recent events in markets within commercial computing, such as Internet access market. For example, at the very outset of the browser-based commercial Internet in 1995, many ISPs wrestled with fundamental decisions about how to commercialize the Internet. Part of the confusion arose from uncertainty about whether to imitate the pricing norms in the bulletin board business, where users phoned into a single server acting as repository for content (hence, the server acted as a electronic equivalent to a bulletin board), or invent a new pricing model.[59] Ultimately, flat-rate pricing emerged in the United States. By 1997, ISPs offered service in every major U.S. city, and many large firms had begun building national networks.[60]

This was but one of numerous experiments to resolve many open questions. For example, a crucial question at the outset concerned the design of the opening page—or, as it was subsequently labeled, *portal*—that users would see when they first clicked on their browser. Should it be a directory of Websites, as if it were the yellow pages for local

[59] This argument is fully developed in Greenstein (2007).

[60] By the fall of 1996, there were over twelve thousand local phone numbers in the United States to call for commercial Internet access, and more than sixty-five thousand by fall 1998. 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.

Websites, or a search tool, permitting a wide variety of tastes, rented from another firm?[61]

The ISPs also varied in the range of services they offered. Different ISPs made distinct choices and learned different lessons about the trade-offs between these choices. No single choice dominated, and as firms learned more, perceptions about the costs and benefits of each changed over time.[62]

While economic experiments give rise to unexpected innovations in computing. Most often the interplay of firms leads to outcomes that none of the firms individually intended. That observation motivates a close examination of learning externalities.

II.9. Economic Experiments with Learning Externalities

Directed experiments are those undertaken by firms for their own purposes, while undirected experiments are those that arise from the interplay of many firms' actions. Learning externalities can arise from both types of experiments and in a variety of ways. Inter*firm* information externalities occur between firms. For example, one firm's directed experiment may teach another firm a lesson, or a set of actions may interact in an undirected experiment and teach every industry participant a lesson. Inter*temporal* externalities, however, occur over time. For example, the lessons of prior experiments may generate lessons on which further experiments are built. In practice, these two externalities are difficult to distinguish from one another.

The positive interfirm information externalities take one of two forms. In one case, what worked for one firm becomes known and imitated by others (e.g., success from an experiment at an ISP in one rural location in 1996 implies it might be profitable in another). Alternatively, what did not work for one firm becomes known and, therefore, avoided: For example the difficulties with the first design for wireless computing became known from experiences in 1997, which caused equipment firms to delay building plans until a more suitable design emerged with institutional support for enforcing interoperability.[63]

---

[61] See the discussion in Haigh (2007).
[62] This can be seen in Greenstein (1999). For a full summary, see Greenstein (2007).
[63] Greenstein (2007) discussed these examples at length.

In another form, one firms' failure can teach lessons that help another succeed. The history of Internet access is littered with examples of failures from which all other firms learned. For example, it is now accepted wisdom that users do not desire only a browser and phone numbers presented as if it were packaged software—as was first marketed by Spry networks in "Internet in a Box." Rather, users want ISPs that offer a different type of service with a different set of market features, combining local services with software tailored to their immediate demands (and tailored to some needs users do not know they even have). It is also accepted wisdom that mass-market users do not desire login names with acronyms that are difficult to recall or do not relate to natural language names, as was widely commercialized by CompuServe. Most users also value avoiding technically laborious set-up costs involving weeks of waiting, as was embedded in early data services, such as Integrated Services Digital Network (ISDN).[64] The list goes on.

Intertemporal externalities also lead to divergence between private costs and benefits and industry-wide costs and benefits. One party (in a directed economic experiment) or several parties (in an undirected economic experiment) assume the cost of generating lessons while many others gain the benefits later. That is, those who pay for lessons in an early market are not necessarily those who use them most profitably in a later market, but no contract between these firms governs the extent of direction of the early investment.

An important feature of intertemporal externalities is the asymmetries to the costs and benefits of generating lessons about commercial failure. Lessons about how to avoid commercial failure can be valuable, but the firm whose failure illustrates the lesson for others rarely, if ever, does so for that purpose, and almost never under contract with the others that (later) gain the benefit of the lessons learned from the failure. In an extreme case, a firm may learn a lesson, teach others from its failure, but go bankrupt before it is able to use that lesson. Even though the lesson was expensive to the stockholders of the firm that initiated the experiment, it was inexpensive to the survivors.

[64] For example, in their estimates of demand for broadband, Savage and Waldman (2005) find that most users are willing to pay a considerable fee to avoid set-up hassles and achieve a reliable service.

Inter-temporal externalities also played a role in the early growth of the Internet. The browser gave many ISPs the confidence to open service for their areas.[65] The growing adoption of the Internet motivated many entrepreneurs to propose new businesses for venture capitalists to fund. The growing adoption of these services by business and households, in turn, motivated other software entrepreneurs to develop business that took advantage of developing electronic commerce, which also motivated further household adoption.[66] In brief, a series of largely uncoordinated, yet complementary, actions by buyers and suppliers reinforced the value of entry and adoption by each other in a positive direction. For a time, that motivated even more adoption and more entry (in the late 1990s) until the ceilings on that value creation became transparent to all, slowing the process. More specifically, after the spring of 2000 adoption by first users continued and increasing use by experienced users also increased, but the rate of entry of new firms declined.

In all these examples, no single firm initiated an economic experiment that altered the state of knowledge about how to best operate equipment or perform a service. Rather, many firms responded to localized user demand, demonstrations of new applications, tangible market experience, vendor reaction to new market situations, and other events that they could not forecast but which yielded useful insights about the most efficient business actions for generating value.

While directed experiments might have partially motivated the actions of any single firm, it would be an error to regard the lessons learned as singularly resulting from only one firm's actions. Instead, the interplay of firms, their actions, and their economic experiments yielded a form of serendipity in learning—learning that resulted from the unanticipated combination of lessons learned from several actions or sources.

II.10. Localization of Innovative Activity in Computing

[65] See the analysis in Downes and Greenstein, 2002 and 2007.

[66] Mowery and Simcoe, 2002a, and Kenney, 2000, contain analyses that place these events in the context of the national innovation system and the structure of Silicon Valley, respectively. See Goldfarb, Hirsh, and Pfarrer (2005) for analysis of the dot-com entry wave in particular.

That innovators learn from one another is no secret to historians of technology. It arises frequently in descriptions of the history of the PC industry, for example.[67] Because learning externalities build over time, however, the accumulation tends to make the externalities geographically localized. In turn, that leads to the concentration of innovative activity in a small number of locations.

There are numerous reasons for this effect. For example, the tacit knowledge about the workings of a prototype cannot be transmitted easily without repeated face-to-face contact, a factor that tends to slow the spreading of information outside of a small geographic region. In addition, the early exploration phases of a new commercial market require giving enormous discretion to entrepreneurial actors or managers within divisions, without formal monitoring mechanisms. Venture capitalists supervising their investments (or managers supervising their employees) may prefer frequent "hands-on" contact, an activity that, once again, is easier in close proximity.

An additional factor shapes the character of a local network of firms, the labor market for technical talent. Firms that share a location (or reside in close proximity to one another) necessarily share a labor market. While the presence of many buyers of labor could bid up wages, the presence of a thick supply of labor also makes it easier to meet unique and potentially short-term demands for specialized skill.[68]

Indeed, the common name for the Santa Clara Valley, the Silicon Valley, recalls the era (1960s and 1970s) when many integrated circuit firms were founded there while sharing a labor market, input supply markets, and financial support structures. This enabled movement of new ideas between firms within close geographic proximity.[69]

---

[67] The early years were distinguished by entrepreneurial energy driving the segment's growth and by firms building innovations on top of other's innovations. See e.g., Cringely (1992), Frieberger and Swaine (1984). Analysts also have emphasized the heavy reliance of third-party vendors on nonproprietary standards, which encouraged vertical disintegration in PC supply. Langlois and Robertson (1992, 1995) analyze the causes and consequences of such vertical integration in great depth for PCs and others facets of electronics.

[68] See, e.g., Fallick, Fleischman, and Rebitzer (2006), and Franco and Filson (2007).

[69] How such areas arise is distinct from questions about how they persist, and I focus on the latter. As for origins of the concentration of the integrated circuit industry, perhaps the best known historical cause was due to the founding of Shockley laboratories, founded by William Shockley in Santa Clara after he left Bell Laboratories, where he, John Bardeen, and Walter Brattain had invested the transistor. For a few years, this was an iconic commercial firm in electronics with a famous founder. With such status, it attracted considerable engineering and technical talent. However, Shockley's overbearing managerial style eventually drove many of the senior managers to leave the firm. Some of them went on to found their own companies, some that later became among the most famous in electronics, such as Fairchild and Intel.

Geographic localization has had several consequences for innovation in computing. For example, early events have had long-lasting consequences for the speed and direction of later innovation. A well-known illustration recounts events at Xerox's Laboratories in Palo Alto, as it was a site that developed key inventions for small-system office computing in the late 1970s. The (then dominant) photocopying company had established a laboratory with many leading researchers in computing. Prototypes for several key designs originated there, including the mouse, the graphical-user interface, and LAN. For a variety of internal reasons, the company was slow to commercialize on these, and, through personnel departures and information leakages, the ideas behind these inventions eventually moved to other nearby companies, such as 3Com and Apple Computer.[70]

Von Burg (2001) argues that sharing information with others was crucial for inviting many firms to make equipment. For example, Bob Metcalfe's design for the Ethernet was built on years of university research in data communications. Metcalfe became a part of Xerox research team. His dissatisfaction with some management decisions led him to initiate his own commercialization effort (3Com), where he found multiple ways to share the technological core of information with others. Subsequently, a community of firms and technologists grew up around the Ethernet standard, and as it became larger others became reassured about its continuation, attracting even more participants. As a result, this community became committed to Metcalfe's design and they collectively enjoyed the benefits of a network effect (pun intended). Eventually, other alternatives could not command much market share, which left the majority in favor of Metcalfe's Ethernet design.[71]

Localization has also shaped the spawning of new businesses. In North America these have tended to be concentrated in a small number of locations, such as the Boston

[70] Apple Computer, for example, initiated programs in graphical user interface and a mouse (famously) after Steve Jobs, then CEO of Apple, received a tour of the Xerox labs.

[71] In fact, a still-active standard-setting IEEE committee 802, which endorsed Metcalfe's design, continues to extend and improve it in directions far outside the scope of Metcalfe's original intent. The IEEE is a nonprofit consortium with representation from the industry. The IEEE also endorsed two competing technologies at the same time. Metcalfe's design, though, eventually became part of a suite of commercial products sold by many firms, including 3Com, which competed against alternative specifications developed by other firms, such as IBM.

area and Silicon Valley.[72] That does not mean all the significant young firms in the last thirty years start in the Valley—after all, for decades the largest large systems computer firm (IBM) was and continues to be head-quartered in New York. Today the largest U.S. PC hardware firm (Dell, founded in mid 1980s) is based in Texas, the largest PC software firm (Microsoft, founded in late 1970s) is based in Washington State, and for many years the largest dial-up national ISP (AOL, founded in mid 1980s) was based in the Virginia/Washington D.C. area. It does mean, however, that many new firms are founded out of or near the Bay Area. In the last thirty years, this includes firms such as Oracle (late 1970s), SUN (early 1980s), 3Com (early 1980s), Cisco (mid 1980s), EBay (mid 1990s), Yahoo (mid 1990s), Google (mid to late 1990s), and many others.

Localized learning displays self-reinforcing features. That is, one successful investment builds on another, with experienced workers and financial institutions continuing to create value. The same area became a nurturing nest for the boom in the PC markets of the late 1970s and early 1980s, the LAN boomlet of the late 1980s, and the dot-com boom of the 1990s. Some analysts have argued that the heavy concentration of invention in Silicon Valley arose from the inordinate extent of sharing of information and mobility among a talented workforce, giving the region greater potential to grow than the greater Boston area.[73] Small firms found the environment nurturing because they could take for granted the availability of many key inputs, a thick labor market for technical talent, financial help from venture capital, and up-to-date information about the latest technical trends.

There are countervailing forces pushing away from the concentration of supply, namely the geographic dispersion of users and the gains to suppliers from collocating next to them. As illustration, Arora and Forman (2006) examine the question of which services are tradable in the outsourcing of IT services. They analyze the outsourcing decisions of a large sample of 99,775 establishments in 2002 and 2004, for two types of IT services: programming and design, and hosting. Programming and design projects require communication of detailed user requirements in contrast to hosting, which requires less coordination between client and service provider. They show that the

[72] There is considerable writing on this topic. See e.g., Saxenian, 1994, Kenney, 2000, Lee, Miller, Hancock, Rowen 2000.
[73] See. e.g., Saxenian (1994), or Kenney and von Burg (1999).

probability of outsourcing programming and design is increasing in the local supply of outsourcing, as should be expected if some non-tradable or "local" component to programming and design services cannot be easily removed. In contrast, the decision to outsource hosting is insensitive to local supply except for users with security concerns.

Remarkably, over time, no single cluster alone has served the entire computer industry. Every facet of the supply chain for computing involves firms headquartered and operating in a much wider set of locations. Entry into facets of these markets has become an important phenomenon world-wide. The supply chain for many complementary components has also been associated with many firms in Western Europe and as well as in China, India, Ireland, Israel, Japan, South Korea, Singapore, and Taiwan. The software industry has spread to different areas of the world.[74] Even more widespread are computing service firms, which follow users dispersed across the globe.

Despite this geographic dispersion over the last five decades, U.S. companies have retained leadership in generating new platforms and commercializing frontier technologies in forms that most users find valuable (Bresnahan and Malerba, 1999). Part of this results from the persistence of platform leadership for a time within a segment. In addition, U.S. firms have historically been ascendant whenever platform leadership has changed. Nevertheless, this pattern seems likely to change in the Twenty-First Century, as non-U.S. firms already have found leadership positions in producing components of many platforms and in related areas of electronics, such as consumer electronics, communication equipment, and specialized software.

The role of localization differs between creation and supply of new computing and its use and adoption. Adoption of frontier applications in computing is not necessarily localized, and on occasion has been much more geographically dispersed. A similar remark holds for the geographic supply for support services for frontier computing technology. As a result, there is little evidence of massively different adoption patterns for new computing across the fifty to one hundred major urban areas in the United States. While data do show slight biases in favor of a few areas, such as San Francisco and Boston, these are readily explained as a result of the composition of the work force or type of industry in the region.

[74] This movement has long antecedents. See e.g., Mowery (1996), or Arora and Gambardella (2005).

More to the point, the differences between major urban areas are small in comparison to the more dramatic differences between major urban areas and slow growing small towns and/or poor rural areas. Those differences show up in PC use, as well as in Internet use and supply, at home and in business.[75]

III. The Commercial Internet in the United States

The networking and Internet revolution in the late 1990s appears to be responsible for blurring once-familiar distinctions between segments of the computing market, geographic locations, and different technologies. For instance, the new networking technologies can build client-server systems within large enterprises and across ownership boundaries. Within these systems are Internet-based computing networks linking potentially vast geographic distances, and thereby supporting the emergence of a "network of networks" on an unprecedented scale.

What, if any, continuity is there in the economic determinants of innovative conduct? The growth of the Internet is a useful case study for illustrating that there is more continuity than meets the eye. Indeed, by identifying such continuity, this review ultimately highlights the very opposite—what is unique in the recent innovative experience.

To be sure, the answer is not obvious. The Internet is not quite like the mainframe or the personal computer. It is not just a single piece of equipment that embodies components from multiple suppliers. Though it helps move data between computers, it is also not quite like the local area network that attaches to existing computing. Its value chain is far more complex and involves many more firms. It is not quite like a new software application. It is not just a single program installed on a computer that generates a new set of functions.

The Internet is also not quite like any commercial communications network that came before it. It is partly a packet-switching network for moving data between computer clients. Yet, that does not fully describe the commercial form it took. A complex

[75] For more on the geography of the Internet in the US, see the reviews in Greenstein (2005), or Greenstein and Prince (2007). For recent evidence on PC use, see Beaudry, Doms and Lewis (2006).

technology had to be embedded in a multilayered network, and many different participants operate its pieces. In addition, the Internet altered many different computer component and software markets simultaneously, so the boundaries of the computing market have changed. A hardware-based definition for the computing market was barely adequate in the 1960s and is no longer sufficient for economic analysis in the Internet era.

Moreover, its commercial diffusion looked quite unlike anything that had come before it. After years of development, a few applications were built that provided compelling value for tens of millions of decisions makers. In the popular imagination this happened overnight, with the creation of the browser. In fact, it had happened over more than two decades, starting from the first government funding at DARPA. The browser was but the last of many innovations, and, thankfully, a commercial marketplace for Internet services had been put in place just before it became available.

A brief review of the size of the Internet access economy gives a sense of how big demand for the Internet became, once it started to commercialize. The enormity of the Internet economy is discussed in Greenstein and McDevitt (2009), which analyzes but a small piece of it, the total revenues for Internet access since the late 1990s. The revenue affiliated with providing access is one of the largest categories of revenue out of the value chain for Internet services, and it is quite large. By 2006 total revenues have reached $39 Billion. That is extraordinary for a technology that had not commercial service providers prior to 1989.

During this growth, the Internet began to accumulate more capabilities and functions, as a range of firms began to use pieces of the Internet to enhance services provided to paying customers. Over time, “the Internet” became a label for not only the Internet but also for all the applications that accumulated around the Internet, used pieces of the Internet, commercialized new functions for the Internet, and which together delivered an enormous array of services to a wide range of users.

Generally speaking, four types of rather different uses share the same capacity: browsing and e-mail, which tend to employ low bandwidth and tolerate delay; video downloading, which can employ high bandwidth and can tolerate some delay; voice-over IP and video-talk, which tend to employ high bandwidth and whose quality declines with delay; and peer-to-peer applications, which tend to use high bandwidth for sustained periods of time, and can tolerate delay, but, in some applications (such as Bit-Torrent) can impose delay on others.[76]

That range of uses and applications serves as cause for both celebration and consternation. The commercial Internet is not just an email network for technically skilled users. It is an email or instant messaging communications network for some, a gaming network for others, a source of news for others, and a distribution channel for video and musical entertainment for others. For many users it is also the principal media for engaging with geographically dispersed communities of friends.

Hence, it is challenging to begin at the basic starting point for empirical analysis—figuring out what to analyze and how to measure its change. This review will use the established frameworks, beginning with known historical facts about technology push prior to the commercialization of the Internet and then moving to a description of the Internet's diffusion and adoption patterns.

III.1. Stretching the Frontier Prior to Commercialization

The Internet began to commercialize around 1992.[77] Within a few years, there was an explosion of commercial investment in Internet infrastructure in the United States. How did that occur? The transfer of the Internet from its research origins to a commercial form involved three somewhat interrelated events: The privatization of the Internet, the creation of the World Wide Web, and the commercialization of the browser. Together these set the stage for a surprising commercial explosion.

A fair reading of the history of each event would suggest that the inventors/initiators did not fully forecast the consequences of their own actions, and most industry insiders were surprised by the changes the commercialization of the Internet

[76] This is explained in considerable detail in Ou (2008).

[77] While the commercialization of the Internet is sometimes dated to the development and implementation of NSF's privatization plan for the NSFNET in 1994-95, that does not recognize the investments made by many early entrants prior to the final NSF plans. Attempts to privatize some assets affiliated with operating the Internet dates to the late 1980s. Those investments resulted in a confrontation between the commercial firms operating the NSFNET under contract with NSF, i.e., IBM and MCI, and those firms, such as Sprint, PSINET, and UUNET, who also were building commercial services and sought to interconnect with other carriers without violating the NSF's Acceptable Use Policy (for traffic that had no direct relationship to research activity). A more proper dating for the beginning of commercialization was ending of this dispute. That occurred with the passage of the Scientific and Advanced Technology Act of 1992, Public Law Number 102-476, sponsored by Rick Boucher (D – 9th District, VA), which amended the NSF acceptable use policy. See Kahin and McConnell (1996) or Hussein (2003).

enabled.[78] That would suggest that a "technology-push" interpretation of the early Internet is consistent with events. While that is partly so, it would be a mistake to go too far with such an interpretation, as if Internet technology was simply dropped on commercial markets like manna from heaven. The early events also cannot be understood apart from the institutional factors shaping commercial behavior at the time.

What became the Internet began in the late 1960s as a research project of the Advanced Research Projects Administration of the United States Department of Defense, the ARPANET. From these origins sprang the building blocks of a new technology for a communications network, one based on sending data where some amount of delay was tolerated. By the mid-1980s, the entire Internet used TCP/IP packet-switching technology to connect most universities and defense contractors.

Management for large parts of the Internet was transferred to the NSF in the mid-1980s. Through NSFNET, the NSF was able to provide connections to its supercomputer centers and a high-speed backbone. Since use of NSFNET was limited to academic and research locations, carriers who carried commercial traffic, such as UUNET, PSINET, and Sprint developed their own private backbones for corporations looking to connect their systems with TCP/IP (Kahn 1995).

By the early 1990s, the NSF had developed a plan to transfer ownership of the Internet out of government hands and into the private sector. The plan for privatization was motivated by several factors. For example, it was forecast (correctly) that a privatized Internet would be more efficient than a government operated one, leading to lower costs for all users. There also was a concern that the NSF could not fund indefinitely the operations of the network, and it was thought that privatization would put the network on more stable financial footing. During the transition another issue arose: several of the private providers of data services were chafing under the NSF's "acceptable use" policy forbidding them to use government-owned assets for commercial purposes. Complete privatization also would remove this issue.

The privatization plan had three important elements. One key element was the operation of data-interchange points. There were precedents for such operations, in

[78] For a long detailed explanation for how unexpected factors shaped outcomes many different Internet markets, see e.g., David (2001) or Greenstein (2007b).

federally operated data-interchange points. Among the earliest private arrangements for this was handled by the Commercial Internet eXchange (CIX), which permitted all parties to exchange data (to "peer" without charge) at locations supported through a group funding effort. NSF's privatization plan led to the opening of several more locations for data-interchange, called Network Access Points (NAPs). Altogether, at the outset these made it possible for multiple actors to enter into the networking business, with no dominant provider able to exclude any other.[79]

A second key element was the privatization of the domain name system, which until that point had been competently handled in a relatively informal operation. Subsequently, it became a large scale thriving (and lucrative) commercial activity. Because it never lost some of its monopoly-like features in its commercial form, the public policies towards the domain system were controversial then and have remained so.[80]

The third and final key element was the shutdown of the NSFNET. When it shut down in 1995, only for-profit organizations were left running the commercial backbone, while a mix of commercial and nonprofit firms operated access points. With the Internet virtually completely privatized, its diffusion path within the United States was largely dependent on market forces and economic incentives.[81]

The NSF privatization plan put in place a scalable network. The operations for updating routing tables, exchanging data, obtaining domain names, and building applications would remain roughly the same even with many more users. Bottlenecks in the provision of capacity also did not materialize because private firms saw opportunities and acted on them, precluding any single firm from operating the network in its entirety and strategically blocking others.

One key early invention arose during privatization, the World Wide Web, and it is linked with a particularly important invention, specifically, the commercial browser. Tim Berners-Lee and Robert Cailliau built key parts of the World Wide Web and, in addition, Berners-Lee organized its pieces into the World Wide Web Consortium that standardized

---

[79] See Hussein (2003) for a review of the changes in these over time.

[80] For different viewpoints, see, e.g., Kesan and Shah (2001), or Mueller (2002).

[81] For a variety of perspectives about the path to privatization, see Abbate (2000), Kahin and McConnell (1997), Hussein (2003), Mowery and Simcoe, (2002a, 2002b), Greenstein (2007b).

many protocols for a growing community of users. Those accomplishments took several years, and at the outset, even Berners-Lee did not forecast their large impact. Indeed, his goals at first were modest and focused on the needs of his employer and a research community he aspired to help.[82] As such, these circumstances fit one factor of Rosenberg's (1995) aforementioned frameworks—a narrow search with results of breadth outside the scope of the initial search.

Specifically, Cailliau and Berners-Lee were employed at CERN, a high energy physics laboratory in Switzerland. They sought to make a program that aided the sharing of textual and nontext files among researchers with a program of sufficient generality to handle many different types of files. Though there had been years of discussion within computer science about how to design such a system, one of Berners-Lee's core insights was *not to* design a perfect system. Rather, he looked for a hyper-text solution that met the needs of the local "constituency," that is, a solution that made it easier for technically-oriented users (physicists who were not computer gurus) to send files easily to one another and make them available for downloading without knowing all the ins and outs of each computer system.

These inventions were useful – but not very useful – unless widely adopted. In 1991 Berners-Lee made two inventions available on shareware sites for free downloading: html (hyper-text markup language) and the URL (universal resource locator) a hyper-text language and labeling system that made transfer of textual and non-textual files easier. Once installed in a host computer these were well suited to Berners-Lee's constituency in two specific senses: (1) It helped users organize transfers of known files and (2) it helped make files available to others without a tremendous amount of pre-transfer searching. Trials and operations over the next several years helped further refine its operations in several ways described below.

As these inventions grew in use, Berners-Lee forecast the need for an organization to assemble and standardize pieces of codes into a broad system of norms for operating in the hyper-text world. He founded the W3C for this purpose. In 1994, he left his employer and established the offices for the W3C in Cambridge, Massachusetts. This organization ultimately helped diffuse many of the software standards and tools that

[82] See, e.g., Berners-Lee and Fischetti (1999), and Gilles and Cailliau (2000).

became important for operating on the commercial Web, such as html and the related tools for deploying it.

III.2. Learning Externalities and Commercialization

Several researchers devised improved versions of browsers that worked on Unix operating systems. These were known among insiders and technically skilled programmers in the research community. Improvements accumulated, and that set the stage for the invention of one additional complement, the commercial browser—an invention that in turn motivated subsequent further inventions.[83]

In 1992, a University of Illinois team situated at the National Center for Super Computing Applications (NCSA), an NSF funded research center that supported a super computer, sought to design an easy-to-use browser for non-researchers. The NCSA supported a large social network of researchers, who regularly used shareware software and made it available to others. Because of this environment (and for other reasons), Mosaic, at first appeared to be a routine project. In this instance, the team of programmers included Marc Andreeson, an undergraduate, and Eric Bina, an employee of NCSA and recent Masters graduate. They took increasing responsibility for the browser project over time, not only to program and design it, but to help it diffuse, to debug it, and to respond to requests for changes from users.

Mosaic had been built on the prior designs, borrowing many elements from them and the programmers tailored the design to the new audience. The project had many features that made it easy to use. After success with their Unix-based project, they developed new features that were novel. They built a browser for a Windows-based systems, at that time the most widely used operating system world-wide for PCs. Up until that point it had not occurred to anyone in the technically adept community of Internet programmers to write something of value for a non-technical user.

The release of Mosaic browser began in early 1993, with the Windows-based browser coming later in the year. It became available on shareware sites aimed at sharing software among university users. Within a year, over a million downloads had occurred.

[83] A version called Viola would shape many of the decisions made by a team who explicitly aimed for a popular browser. See Gilles and Caillaiu (2000).

As it grew in popularity, the managers at the University of Illinois, Urbana/Champaign realized this invention had great potential. They arranged for commercial licensing of the browser. They anticipated that the browser would diffuse into popular use through both shareware, which was free, and commercial licensing, which involved a more organized commercial support. Almost certainly those expectations would have been correct had a third channel not emerged. That third channel involved the student programmers—but none of the faculty or other administrators—who had helped develop the Mosaic browser; the students decided to start a business around the same time as the university began its licensing program.

Specifically, Marc Andreessen, one of the lead programmers on the Mosaic project, had graduated in December of 1993, moved to the west coast to take a software job, and subsequently grabbed the attention of Jim Clark, founder of Silicon-Graphics from several years earlier. Clark had excellent established connections with the West Coast computing community. Clark's and Andreesen's relationship coalesced into a business plan in the spring of 1994. They called themselves the *Mosaic Communications Company* and sketched a plan to make money selling a browser. They received venture funding from the same venture capitalists that had backed Clark's earlier efforts, hired many of the same programmers who had worked at NCSA in Champaign, and went on a crash course to become a large organization supporting worldwide use of their browser. Eventually this plan would expand far beyond the browser, blossoming into an extensive business plan to support a range of complementary activities around their own browser, server tools, and range of services. In effect, this program eventually aimed to make the licensing program of the university obsolete.[84]

The University's officially sanctioned channel was managed by a third party—a company known as Spyglass, located in Illinois with a history of commercializing NCSA inventions. Spyglass was given the right to license the trademarked name *Mosaic*. Spyglass eventually decided to defend its intellectual property, forcing Mosaic Communications Company not only to change their name to *Netscape*. The threat of further legal problems also made Netscape's programmers take extra care not to overlap

[84] See, e.g., Cusumano and Yoffie (1998).

with the intellectual property owned by the University.[85]

Netscape's beta browser was released in the late fall of 1994, gaining publicity, followed by the first commercial release in the winter of 1995, and its IPO (initial public offering) in August of 1995. Netscape's business model and marketing efforts were wildly more successful than any other licensee's at catalyzing many other market participations. Among the results from its many effects, Netscape's activities were instrumental in motivating Bill Gates to reverse his previous position about staying out of the browser business.[86]

Indeed, the University of Illinois played a role there too. After failing to buy Netscape in the summer of 1995 or deter them from pursuing a strategy that conflicted with Microsoft's desire to lead all software application development, Microsoft set on a course to offer a browser. The fastest way to do that was through licensing Spyglass' browser, which had been licensed by farsighted Microsoft employees in January of 1995 (as part of an internal campaign to change the priorities of the organization). After doing so, Microsoft added a few features, re-branded the browser, and unveiled Internet Explorer in December 1995.[87]

The University of Illinois also served as the origin of the most widely used HTTP (Hypertext Transfer Protocol) server, a collection of technologies that supported browsing, particularly SGI script, and use of web technologies. That eventually became known as the Apache server. The emergence of Apache was not the original intent of the university's administrators. The server was available for use as shareware, and many Web masters took advantage of it, adding improvements as needed.

By February 1995, however, the university was not keeping up with users, partially because key employees had left for Netscape. Many users of the NCSA server software sought to coordinate further improvements. NCSA tried to revive the software

[85] Netscape's management claimed it would have reprogrammed the browser from the ground up in any event, because they were developing software to support their long run goals, which required starting from scratch. However, the concerns about intellectual property made that goal a necessity rather than a luxury.
[86] See Bresnahan, Greenstein and Henderson (2007). For the original memo by Gate announcing this change in direction, see Gates (1995).
[87] See the account in Sink (2007). He argues that Spyglass' management was quite elated to get Microsoft as a licensee. After they licensed the browser Microsoft increasingly devoted more programmers to it over time, positioning itself as the firm to support other application development. All the other licensees eventually chose to either get support from Microsoft or Netscape; none continued with Spyglass.

in April 1995, but, upon learning about the Apache effort, changed its plans to support a shareware version of the server software, and cooperated with Apache.[88] Out of these efforts arose the most widely used server on the commercial Internet and it became one of the most widely adopted open source projects after Linux.

Both the browser and the server are good examples of the role of university-funded research breaking new technological ground that commercial firms overlooked. It is also good illustration of the value-added that commercial firms bring to products, refining them, branding them, and servicing them, contributing to the diffusion of a technology in a myriad number of ways.

It is also an example of localization in the commercialization of software production. Neither browser or server technology moved to firms headquartered in Illinois, despite the seeding of a Spyglass, or the nearby presence of Chicago, the third largest city in the country and one that does not lack the appropriate infrastructure or labor market for technical talent. To be sure, production did stay inside the United States. Clark was already located in the Bay Area, and there was no debate that the new firm would stay there to take advantage of the existing cluster of software firms, venture capital finance, and labor market for technical talent. Their strategy depending on scaling their organization rapidly through hiring talent – arranging many business ventures and growing with the help of many complementary firms – and the location was a key part of that.[89]

Accordingly, while Netscape was the most important browser firm in the Internet, the core of development took place in California. After a few years, Microsoft's Internet Explorer became more important than Netscape, and for that period the key production for the browser was located just outside Seattle, where Microsoft located what became the Internet product and tools division.

As an interesting contrast, Apache was founded by several people and some of the

[88] The name apache was a play on words, as the first February, 1995, effort involved bringing together a piece of software that involved many "patches." See http://httpd.apache.org/ABOUT_APACHE.html, accessed March 2007.

[89] The reliance on external complementors is described in detail in Cusumano and Yoffie (1998). While they do not highlight the role of geographic localization in that strategy, it was plainly obvious to everyone that the users were geographically dispersed, many of the complementors would be geographically close, and most of the software vendors had an affinity for Netscape's competitive goals in opposition to Microsoft.

key leaders were located in (and remain in) the San Francisco Bay Area. In some sense, however, it is not located anywhere in particular, remaining largely a virtual organization, with member and code coming from programmers all over the world.

III.3. Localization and Commercialization of Internet Software Production

As the browsers began to diffuse, entrants began exploring new businesses. Many of them were founded in the San Francisco Bay Area, a direct outgrowth of the entrepreneurial-orientation of the venture-funded community, who was looking to fund start-ups, and the pre-existing labor market for technically skilled employees at other firms, who were becoming privy to the latest developments for the commercial web. These entrants included Yahoo!, Hotmail, Excite, EBay, Vermeer, and many others that anticipated building businesses on browser-based computing.[90]

The entrepreneurial movement at the time extended far beyond the San Francisco boundaries. For example, Bill Gates' memo of May 1995 entitled "The Internet Tidal Wave" advises other executives to examine Web pages by Lycos, Yahoo!, Oracle, Symantec, Borland, Adobe, Lotus, the *San Jose Mercury News*, Novell, Real Audio, Disney, Paramount, MCI, Sony, ESPN, and many other sites.[91] Although the first eight of these companies were located in the Bay Area, the rest were not.[92] As another example, in May of 1995, *Boardwatch* magazine, the primary trade publication for the U.S. bulletin board service market, listed over 700 price plans for getting commercial Internet access from local ISPs all over the United States.[93]

These types of events fueled expectations among industry insiders, futurists, and venture investors that substantial demand for the Internet at households and businesses

[90] With twenty-twenty hindsight it is too easy to treat this outcome as obvious. For insights into the uncertain experience of an entrepreneur in this early era, see Ferguson (1999).

[91] See Gates (1995). This memo announced Gate's intention to alter the strategy direction of the company and focus on Internet technologies. The list sites were there to give other executives examples of sites Gates found exemplar in various respects.

[92] As yet another illustration of the extent of geographic dispersion consider this example from a slightly later time, the analysis of commercial prospects in the Internet market – first performed in late 1995 by Morgan Stanley analysts Mary Meeker and Chris DuPuy (1996) after Morgan Stanley successfully handled Netscape's 1995 IPO. It highlights young Internet companies from California, such as Silicon Graphics, SUN, Cisco, Excite, Netcom, and Netscape, but also features plenty from all over the country, such as AOL, CompuServe, UUNet, Dell, Compaq, US Robotics, IBM, and others.

[93] This is a conservative estimate. Many more firms were list, but many did not provide pricing information. See Stranger and Greenstein (2007).

would emerge quickly.[94] Many firms had began to initiate projects for converting part of their business to browser-based computing, especially among those technology firms whose livelihoods depending on racing into mass markets faster than others. The Netscape IPO took place in August of 1995 and it was wildly oversubscribed (or very badly underpriced), pushing the trading price multiple times higher on the initial day of trading.

With the Internet, the relationship between the investor community and entrepreneurial community took a different scale and pace than it had in prior technology-induced waves, such as with PCs, LANs, and client-server systems. In part, this was due to the breadth of perceived opportunities. Rather than being a brief race among several dozen firms to develop new components and related systems, the Internet invited a wide range of new thinking across many activities—in back-office computing, home computing, and information retrieval activities in numerous information-intensive industries, such as finance, warehousing logistics, news, entertainment, and more. Ultimately, the Internet motivated the entry of new entrepreneurial firms continuously until the spring of 2000, peaking in 1998 and 1999. The entrants ranged over the entire spectrum of businesses shaped by the emergence of IT in the prior decades.

A new data carrier industry also arose. It involved a mix of existing firms and new ones, especially among ISPs. More than 5000 such firms arose in the United States by 1998, most of them small and specialized on serving a local area.[95] A few prominent firms emerged, principally AOL, Earthlink, Mindspring, Netzero, Level3, PSINET and a few others. They competed with new divisions at established firms, such as AT&T and MCI.

Perhaps the most attention went to a category of entrant, generically known as the *dot-com.* The label came from a feature of the domain name system for the Internet, which initially designated five types of domain names: gov, net, org, edu, and com.[96] *Gov* was for government entities and *edu* for educational institutions. *Net, org* and *com* were

---

[94] For different perspectives on how U.S. business reacted to the new opportunity, see, e.g., Mowery and Simcoe, 2002a, 2002b, Kenney 2000, Greenstein 2001.

[95] For detail, see Stranger and Greenstein (2007).

[96] Every country was assigned control over the allocation of domain names underneath its two-letter country code. The country code for the United States is US, but com, net, edu, gov and org presume US sovereignty. See Mueller (2002) for a description of how this developed.

designated for non-profit and private entities, organizations, and networks. *Com* became the most popular among commercial firms (by far), even for firms not based in the United States.

Dot-com entrants covered a wide range of new businesses. Initially a large number of entrants went into advertising-supported Websites or sites with no usage fees.[97] A number of sites covered directory and search activities, while others specialized in supporting conversations and information sharing for particular topics or groups.[98]

An entirely distinct group went into electronic commerce in various forms. Some sought to sell and distribute goods and services, such as books or travel services or apparel. Others sought to sell subscriptions to services, such as the *New York Times* crossword puzzle or regular updates of industry news. Still others tried to assemble groups of buyers and sellers, either in open auctions or for more tailored matching purposes.

At the time, there was a pervasive ideology about using the Internet to remake economic activities. Given the label *the new economy* in popular discussion, this vision argued for the Internet's exceptional nature. It argued that the Internet enabled business processes which would remake the structure of production and the redesign the locus of decision making within organizations. Its enthusiasts encouraged entry by new firms on the premise that many incumbent firms would be unable to adapt the new technologies to their existing businesses fast enough or effectively enough to compete with the new entrants.[99]

Despite such rhetoric, existing firms did not stand still. Many leading firms invested in new processes and new Websites to support their businesses. Still others hired consultants and experts to help with adjusting to the new potential and new commercial threats. Accordingly, a large consulting and advising industry grew in the late 1990s at companies such as Accenture and EDS and many others, as did a substantial set of firms

[97] See Goldfarb (2004).

[98] See Haigh (2007) for a discussion of this segment, and see Goldfarb, Hirsch and Pfarrer (2005) for discussion of the variety of entrants.

[99] Some of this view was founded on serious academic research, e.g., Bower and Christensen (1997) or Christensen (1997), but over time it developed into an untested belief system that did not necessarily measure itself against results. While it justified numerous IT projects, ultimately these were judged against their productivity (e.g., see the summary in McKinsey, 2001). For a variety of perspectives, see e.g., Forman and Goldfarb (2006), Goldfarb, Kirsch, and Pfarrer (2005), or Shiller (2000).

for supply software infrastructure, such as hosting and Web design.

Although the entry and investment boom around the Internet resembled prior entrepreneurially led booms in new technical opportunities, such with PCs, LANS, and client-server networks, it also differed in size, scope, and aspiration. With the Internet some new entrants sought to create new businesses and remake value chains from production to distribution to final user. There was open discussion about changing the entire chain of actions supporting the delivery of any valuable information-intensive activity, such as music, publishing, news, financial activities, and entertainment.

For many investors a pervasive optimism in the late 1990s moved the otherwise enormous uncertainty about the source of value into the background. Nonetheless, there were open questions about how large the advertising-supported economy would become in its on-line format and how quickly on-line retail channels would become used by buyers and whether their growth would favor existing firms or incumbents. There were also questions about how to design for the most valuable features, computing languages, and features to serve emerging usage patterns among the growing set of Internet users. In brief, this period was marked by economic experiments across a wide range of activities that overlapped with applications of computing, as well as any up-stream or down-stream activity related to it.

III.4. Platform Development for the Commercial Internet

As one might expect in a market with divided technical leadership, commercializing the Internet introduced a plethora of variants on developing platforms. There was considerable disagreement among participants about whether any of these were valuable, and about what to call each distinct strategy. Even with hindsight, there is disagreement about which strategic investments worked and why.

One approach to new platforms employed existing assets as much as possible, by trying to encourage existing users to employ a standard bundle of components that continued to use the proprietary assets of an existing firm. Firms with leading platform positions in the computing or data-carrier business before the Internet commercialized, such as IBM, MCI, Microsoft, Intel, Novel, 3Com and Cisco, initially sought ways to create value for their users through incremental innovation, while simultaneously attempting not to lose their leading positions. Beyond that generality, the details differed significantly

I now explore some of the dramatic events arising from the actions that individual firms and types of firms undertook to develop successful platforms. I begin with IBM: At the outset of the 1990s, before the Internet commercialized, IBM's mainframe business had begun to decline significantly. That led the board to remove the CEO (chief executive officer) and break with precedent by hiring a CEO from outside the company, Louis Gerstner. Along with selling off several business units such as the networking carrier business, Gerstner chose to concentrate the firm's assets in advising companies how to implement IT in effective ways that created value. This approach used IBM's existing relationships with enterprises and grew the company into essentially a large consulting and systems integration practice, which grew even larger through mergers with smaller firms.[100]

Cisco behavior also drew attention in popular discussion for a few distinct reasons. It sought to become the leading provider of enterprise-level data equipment, through internal growth and designs out of its own considerable Research and Development department. Notably, though, Cisco also sought to increase its product line and set of personnel through the acquisition of small venture-funded start-ups. Most of these had one prototype product, if that, and fewer than 100 employees. Over the course of the late 1990s, Cisco bought over eighty such companies, forming an ecosystem with the venture capital community who started those firms.[101]

To be sure, venture-funded entry was not new for computing markets. The new phenomenon was the systematic strategy of larger firms to grow their portfolio of innovative projects through venture funded firms. Firms, such as Cisco and JDS Uniphase, would monitor external innovative projects and "cherry pick" those that fit with their strategies, either through merger/acquisition or other cooperative forms of cooperation.[102]

---

[100] This account oversimplifies a enormous task. For more detail, see Gerstner (2002) or Carr (1999).

[101] There was no secret about what Cisco did, but the number of deals and their frequency was without precedent. See, e.g., Paulsen (2001).

[102] For more on these type of strategies, see e.g., Paulsen (2001), Arora, Fosfuri and Gambardella (2001), Gans and Stern (2002), Gans, Hsu and Stern (2003), and Gawer and Cusumano (2001).

Among all incumbent computing firms who responded to the opportunities in the commercial Internet, Microsoft's actions perhaps gained the most attention in popular discussion because it generated a series of events known as the browser wars, which involved a dramatic confrontation between Netscape's browser business and Microsoft's. In addition, Microsoft's aggressive competitive tactics motivated a large federal antitrust case against it.

At a basic level, there was no mystery at all to Microsoft's goals. The company had gained an extraordinarily lucrative position in the PC operating system and applications markets, and a growing position in the lucrative networking operating systems market, and it had an incentive to protect that position by preventing any alternative platform from emerging, if it could.[103] One step to do that involved providing a browser, eventually named Internet Explorer. A second step involved converting most application developers to making programs compatible with that browser and no other. The second step, of course, reduced the value of Netscape's platform, but could not be achieved without first persuading users to employ Internet Explorer for most of their browsing, thereby giving Microsoft increasing ability to bargain with developers. Indeed, eventually Microsoft did achieve its first goal, making Internet Explorer the default browser on most PCs. Eventually they gained share of Internet Explorer that, for a time, exceeded 90 percent of all Internet users.[104]

If the goals, by themselves, were not controversial, the tactics were. Late in making crucial investments in Internet technologies, Microsoft was technically far behind its nearest rivals once it finally did begin to make investments to support its proprietary platform. Then, concerned that others would develop a persistent platform, Microsoft used its existing relationships with firms to insist on contractual obligations (such as first-screen restrictions, "de facto" exclusive deals, and controversial quid-pro-quos) that restricted innovative actions of OEMs. Those tactics came under close scrutiny, and once questioned, Microsoft reacted with an aggressive legal response instead of finding any

---

[103] Gates (1995) provides a quite cogent analysis as to why the Internet could give rise to platforms that would undercut Microsoft's margins. In brief, a browser controlled by another firm could support a platform comprised of software available on the Internet that eliminated the unique position of Microsoft's operating system, reducing its value.
[104] See e.g., Cusumano and Yoffie (2000), and Bresnahan and Yin (2006).

point of compromise. That confrontation played a role in inducing the Department of Justice to bring a federal antitrust suit, which kept embarrassing revelations about the browser wars in the news long after Microsoft had achieved its commercial goals.[105]

Intel, like the previous examples, also altered its innovative priorities for computing in response to the new opportunities. Intel had crept into the motherboard business over the 1990s as it initiated a variety of improvements to the designs of computers using its microprocessors.[106] It also accelerated investments in manufacturing technology, resulting in what appeared to be the fastest rate of improvement in price/performance ever achieved.[107]

Another series of entrants sought to take advantage of the new opportunities to create value with Internet technologies: A large variety of different firms sought to match buyers with sellers, or match users with similar communities of interest, or even match questioner and seeker with informative answers. Economists have labeled these different approaches either *intermediation* (Lucking-Reily and Spulber, 2001) or *two-sided markets* (Evans, Hagiu and Schmalensee, 2006). These initiatives varied in their pricing structures, that is, in how they sought to generate monetary revenues from the new services offered. Some firms sought to make revenue from charging subscription fees (e.g., dating services), while others sought to make it through advertising next to the relevant information.

A related important innovative offering were *portals.* Portals became some of the most popular Websites among Internet users, accounting for the opening pages for most users, acting as a tool for organizing web surfing, and accounting for the highest share of an individual's time on line. There were two approaches to providing portal services at first. One approach made it an extension of another service, and virtually all of these firms aimed at mainstream users. For example, Netscape and Microsoft made the opening Web page a default setting on their browser, directing considerable traffic to

[105] For different accounts and viewpoints, see, e.g., Cusumano and Yoffie (2000), Henderson (2000), Fisher (2001), or Bresnahan (2004), Schmalensee (2000), or for a sardonic take with biting insight, see Chapman (2006), chapter 10.
[106] See Cusumano and Gawer (2002), and Gawer and Henderson (2006).
[107] See e.g., Aizcorbe (2006) for analysis of the determinant of improvements, as well as Flamm (2003). See Aizcorbe, Flamm and Khursid (2007) for why this helped computing equipment more than communications equipment.

netscape.com and msn.com, respectively. Both were far down in use to AOL, which made its opening page an extension of its dial-up Internet service and organized its contents in an easy-to-use proprietary format—a strategy that became known as a *walled garden.*

Another approach was of stand-alone portals, and these differed in their appeal to technical and non-technical users. Most early (and successful) portals in the mid-to-late 1990s provided directory services for the vast amount of content on the Web. Yahoo!, Lycos, Excite, and others took this approach. Yahoo! grew into a vast organizer of content and a popular destination. Excite sold itself to @home as part of a strategy to help cable Internet users. Late in the 1990s, a newly founded firm, Google, entered with a new approach to searching—ranking Web pages on the basis of the number of links to them, and supporting itself with advertising. This turned out to be popular with users; after the turn of the millennium, Google continued to grow this business and eventually surpassed the early leader, Yahoo!, in usage, revenue, and capitalized value.

This brief survey is the tip of the iceberg of how innovative computing enabled numerous new opportunities in electronic commerce and information-intensive activities, linking the market structure of computing and many other firms.[108] Retailing firms began developing strategies for differentiating their services from each other, and for customizing distinct versions of their activities to many customers. Firms in media markets – news organizations, music distributors, entertainment companies – began developing software solutions to restructure their services for the Internet. In turn, these opportunities generated a cornucopia of innovative efforts at software and hardware computing firms to support this new direction for value creation.[109]

Unlike the most Apocalyptic predictions after the episode in browsing, no dominant and proprietary Internet platform emerged after the first decade of the commercial Internet with control over the entire commercial value chain. That is, while

[108] Other prominent platforms include those provided by Research in Motion (Blackberry), Apple (iPhone, iPod), Oracle (enterprise databases), E-Bay (auctions), Facebook (social relationships), as well as many others. Each one of these platforms deserves a longer description, and the reader should be clear that the absence of that here is due to space-constraints, not their lack of importance.

[109] For more on the economic factors shaping competitive activity electronic commerce and technically-intensive information industries, see Shapiro and Varian (1998), Varian (2000), Smith, Bailey and Brynjolfsson (2000), Bakos and Brynjolfsson (2001), Bakos, Brynjolfsson and Lichtman (2002). For an overview of market structure, see Varian (2001).

many firms succeeded in making considerable revenue during this period of growth and in controlling niches or corners of the value chain, there was no single proprietary firm largely shaping the direction of technical change in all aspects. And, yet, at the same time, something familiar was emerging: In each subcomponent or horizontal layer of the internet, a few prominent firms made viable businesses selling components and services that the vast majority of users continued to use. This was a step towards becoming a standard bundle for what would someday be recognized as an Internet platform.

This observation can be restated as an aphorism. Since it first began to diffuse to the general public, the Internet has been called a "network of networks." Yet that phrase is misleading. It does not reflect how commercial behavior shaped the evolution of technology in the last decade and a half. Leading firms and their business partners view the commercial Internet through the same lens they view activities in the rest of computing. For them, the commercial Internet is a "network of platforms."

III.5. New Forms of Organizations to Coordinate Firm Conduct

Some observers attributed the rapid accumulation of experimentation to the emergence of a new form of leadership for designing standards, one that involved collections of market participants. The standards committees that were responsible for designing key standards for the Internet were comprised of representatives from many firms and interested researchers from universities and other non-profit organizations. Because undirected economic experiments are those undertaken by more than one firm working together, by definition, the committees participated in these types of experiments. This raised the profile of activities inside standards committees and it directed attention at different forms of consensus-oriented standards processes for designing standards accommodating a variety of complementary goods and services.

Ultimately, the accumulation of Internet industry knowledge depended on spreading the lessons learned from economic experiments. Further innovations then built on that knowledge, renewing a cycle of accumulated lessons from more experiments. This accumulation was a key driver of the market's evolution because it set the conditions for innovative behavior. Standards committees participated in this cycle and helped shape the Internet by affecting, for example, pricing, the quality of services, and the identity of leading firms.

Standards committees had always played some role in the computer market. Their role in the Internet was more notable for what it was *not:* These institutions were not beholden to the managerial auspices of AT&T or IBM. For that matter, these committees also did not simply ratify the design decisions of Intel, Microsoft, or Cisco, though all those firms sent representatives who had a voice in shaping outcomes.

The range of such important decisions shaped by standards committee was without precedent. The IEEE, for example, made designs that shaped the LAN market, modem, and wireless data communications markets, while the IETF made designs that shaped the operations of every piece of equipment using TCP/IP standards.[110] Many of these decisions went into use quickly, ensured that all complying components would interoperate, and had enormous consequences for the proprietary interests of firms. Never before had such a large industry had so much of its innovative activity shaped by collective firm decisions.

Another notable feature of these committees was their governance structure. They were largely organized without much government directive or mandate, especially at the outset of the Internet. It is incorrect to say that the government was uninvolved: After all, the NSF and Department of Defense both had played a crucial role in starting and sponsoring organizations that managed and improved the operations of the Internet.[111] Often, government representatives were present and influential on specific features of design, and sometimes government provided crucial endorsements. It is just that at the outset of commercialization, governments did not put the force of law behind many of these standards, or insist that government employees have exclusive influence on design choices. Rather, the commercial Internet of the 1990s embodied the accumulation of multiple improvements suggested through a process of consensus in committees, and that consensus depended in large part on the functional ability of code to perform as

---

[110] Simcoe, 2007, provides an overview of the operations of the standardization process at IETF and its changes as it grew.

[111] This is especially true of the Internet Architecture Board and Internet Engineering Task Force, before it moved under the auspices of the Internet Society in 1992, where it remains today. See e.g., Abbate (1999), or Russell (2006).

claimed.[112]

Its importance can be understood in comparison to the next closest alternative. The next closest alternative design for global networking— the Open Systems Interconnection model, a.k.a., OSI seven-layer model—arose in the late 1980s in a process almost as different as could be any far-reaching standard. The OSI was a formal standard design for interconnecting networks that arose from an international standards body, reflecting the representation of multiple countries and participants. The network engineering community in the United States preferred their bottom-up approach to the OSI top-down approach, and, when given the opportunity, invested according to their preferences.[113] Indeed, pieces of the OSI model still live on today (how much is a point of some debate), but the locus of decision making over the direction of Internet standards resides firmly outside of what's left of the organization.

The lack of government involvement could also be seen in other aspects of the Internet in the United States. For example, although the Federal Communications Commission (FCC) had mandated a standard for digital television (as it had for color television), it refrained from mandating most Internet equipment design decisions. Just as the FCC had not mandated Ethernet design standards, so it let the spectrum become available for experiments by multiple groups who competed for *wireless* Ethernet standards, which eventually became Wi-Fi. Similarly, the FCC did not mandate a standard for modems other than to impose requirements that limited interference. It also did not mandate interconnection regulatory regime for Internet carriers in the 1990s, explicitly letting the firms innovate in the structure of their business dealings with one another, and evolve those dealings as they saw fit.[114]

Another notable innovative format for consensus decision-making went by the name, *Open Source.* These organizations used variants on the General Public License

[112] See Abbate (1999) for a history of the design of these protocols. See Partridge (2008) for a history of the processes that led to the development of email, for example.

[113] See e.g., Russell (2006).

[114] The latter forbearance was deliberate. On the lack of interference in the design of the Ethernet, see von Burg (2001). On the design of 56K modems, see Augereau, Greenstein and Rysman (2007). On the lack of regulation for network interconnection, see the full discussions in e.g., Oxman (1999) or Kende (2000) or the account in Neuchterlein and Weiser (2005). More recent experience has departed from these trends, particularly in principles for regulating last-mile infrastructure. A summary of these departures is in Greenstein (2007a).

(GPL), created by Richard Stallman. He called his creation *copy left* and he intended it as a contrast to the use of copyright. A GPL required all contributors to give up ownership rights to the collective group. Any contributor could use any others' code and build on it, but only if no single individual and no firm claimed rights to exclude others from using their contribution.

The use of the GPL in the 1990s differed from Stallman's original goal. He had wanted to create a format that reduced the proprietary discretion restricting a user's ability to change computer code. By the mid-1990s and beyond, many Open Source licenses employed many variants that sought to accommodate commercial activity in one form or another, while still allowing users to contribute code and share with one another.

One well-known Open Source project was Linux, a basis for computer operating systems. It was begun by Linus Torvald in the early 1990s as a derivative, or "fix," to Unix. It was freely distributed, with alternating releases of a "beta" and "final" version. Starting around 1994 to 1995, about the same time as the commercialization of the Internet, Linux began to take off. What had started as a small project caught the attention of many Unix users both at commercial companies and elsewhere. Many users began converting from proprietary versions of Unix (often sold by the hardware manufacturers) and began basing their operating systems on Linux, which was not proprietary. Many of these same users began contributing back to the Linux project, strengthening the range of applications. This movement gained so much momentum that Linux-based systems became the most common server software other than Microsoft software. Many firms, such as Red Hat, began to turn a profit selling services and related components.[115]

As noted earlier, Apache was another early project founded to support and create "fixes" for the HTTP Web server originally written by programmers at the NCSA. By 2006, more than 65% of Websites in the world were powered by the Apache HTTP Web Server.[116] Apache differed from many other Open Source organizations in that contributors "earned" the right to access the code. To be a contributor one had to be

---

[115] There is considerable writing about the growth of the production of Linux software and from a variety of perspectives. See. e.g., Dalle, David, Ghosh, and Wolak (2004), Lerner and Tirole (2002), Von-Hippel (2005), West and Gallagher (2006), or Arora and Farasat (2007). For an account and analysis of how many firms got on the Linux bandwagon, see, e.g., Dedrick and West (2001), or Fosfuri, Giarranta and Luzzi (2005).

[116] For more on the history and operation of Apache, see, e.g., Mockus, Fielding and Herbsleb (2005).

working on at least one of Apache's projects. By 2006, the organization had an active and large contributor base and seemed poised to continue indefinitely.

Another early Open Source project, MySQL, pursued a distinct model. From the outset, the organization behind MySQL aspired to make a profit. MySQL was a database with Website-powering, packaged software, and enterprise applications. Many small to medium-sized companies utilized the free basic MySQL for their operations. This package was developed through an Open Source arrangement, in which no user paid for use of the software. Millions of copies have been distributed and companies pay only for the more advanced features of MySQL. The number of paying customers numbers in the tens of thousands; and the company only makes its revenue from these customers.

Perhaps the most well-known Open Source format was the least technical. It originated from something called a *wiki*. Developed in 1995 by Ward Cunningham, a software engineer from Portland, Oregon, wikis can either be used in a closed work group or used by everyone on the Internet. They originally were formed to replicate or make a variation on existing products or services, with the purpose of fixing bugs within the various systems. Accordingly, wikis were first developed and intended for software development, but had grown out of that first use and were now applied to a multitude of applications.

A particular popular application of wikis, Wikipedia, garnered world-wide attention. In the case of Wikipedia, the format was applied to the development of textual and nontextual content displayed on the Web. Wikipedia beat out Microsoft's Encarta for the honor of the Internet's top research site in 2005, a position that it has held thereafter, and it has been consistently ranked as a top twenty site for all Internet users. It is an online-only encyclopedia. The content is user-created and edited. As its homepage proudly states, it is "The Free Encyclopedia That Anyone Can Edit." The site has always been free of charge and never accepted advertising.[117]

These new organizational forms gave rise to a different leadership structure with regard to innovations for computing and the Internet – or what some observers regarded as the absence of a leadership structure. Unlike prior episodes of new uses for computing, after the first decade of the commercial Internet, no single firm or small

[117] For more information, see Greenstein (2006).

group of firms emerged as the key drivers of technical change. Technical leadership did play some part in this behavior, but the additional and numerous Open Source initiatives across a range of applications were a driving force that led to this uncoordinated innovative behavior.

These experiments in new organizational forms gave rise to the belief, quite commonly expressed during the late 1990s, that the "new economy" would alter not only the use of computing, but also its production. It became apparent to savvy observers, however, that the extreme version of that prediction would not come to pass. Proprietary software continued to flourish, often co-existing with Open Source software while competing for some boundary uses. In fact, in many cases, such as Linux, the commercial firms actually prospered as part of the large coalition of firms supporting the Open Source software. In the case of MySQL, it was eventually purchased by SAP.

Open questions festered, though: What types of innovative activity would remain proprietary and what applications were best organized around Open Source software? More broadly, if the fundamental economics of platforms still held in spite of divided leadership and the importance of nonproprietary standards development, then what form will new innovations take? In addition, how did focal developments arise in the absence of proprietary interest in sponsoring platforms?

These events illustrate several different mechanisms for the emergence of focal points for development other than sponsorship by a firm. For example, an individual with technical skill, a history of inventiveness, and, perhaps, celebrity status can gain authority over a range of standards, as Linus Torvald or Tim Berners-Lee did. Alternatively, a standardization institution with visionary leadership can take such a role, as several of the IEEE committees did for wireless Ethernet standards, and as the IETF did for a range of technologies in networking using the TCP/IP stack.

Moreover, both users and vendors have incentives to see focal points continue. So, as with any platform, these coalitions of supporting firms are difficult to stop once they gain a self-sustaining size. In brief, vendors continue to invest in these coalitions, as long as the hardware, software, and services embedded with these standards continue to provide functionality that participants find valuable.

III.6. Diffusion of the Commercial Internet

From the previous sections, it is clear that like computing, the Internet was associated with high fixed invention costs and low marginal re-use costs. It also generated heavy early investment and led to frequent repurposing of focal inventions. In addition, the diffusion of the Internet followed the predictable regularities of a GPT. For example, it always takes time to move a frontier technology from a small cadre of enthusiastic first users to a larger majority of potential users. In this sense, the economic patterns found throughout the early diffusion of the Internet are general.

Nevertheless, the diffusion of the Internet also possessed some unique features. It thus far proceeded in two waves. There was a clear difference between low-speed/dial-up connection and high-speed/hard-wire connection. In the early 1990s, those with dial-up connection were considered at the frontier, but by the turn of the millennium, at households the new frontier consisted of high-speed connections, mainly through (Digital Subscriber Line) DSL and cable modem supported access.

According to the National Telecommunications Information Administration (NTIA) (2004) study, as of 2003, approximately 61.8% of American homes owned a PC, with Internet participation rates at 54.6%. By 2006 Internet participation had reached 73% of households, with only a declining minority using dial-up connections.[118] These adoption rates suggest that the diffusion of each technology is moving into the late majority category of adopters, though there is considerable disagreement about how to portray the rate of adoption for the remaining households. Any entity (household, individual, or firm) is considered connected to the Internet if it has the capability of exchanging information with other entities via the physical structure of the Internet. Connections come at different speeds (56K dial-up versus broadband) and from different types of suppliers (AOL versus a telephone company). Figure 1 provides a visual representation of this diffusion to households, as well as the transition into broadband.

Data from 2001 show Internet usage to be positively correlated with household income, employment status, and educational attainment (NTIA, 2002). With regard to age, the highest participation rates were among teenagers, while Americans in their prime working ages (20-50 years of age) were also well connected (about 70%) (NTIA 2002).

[118] See e.g., Greenstein and McDevitt (2009).

Although there did not appear to be a gender gap in Internet usage, there did appear to be a significant gap in usage between two widely defined racial groups: (1) whites, Asian Americans, Pacific Islanders (approximately 70%); and (2) Blacks and Hispanics (less than 40%) (NTIA 2002). Much of this disparity in Internet usage can be attributed to observable differences in education and income. For example, at the highest levels of income and education there were no significant differences in adoption and use across ethnicities.

Since the vast majority (87.6%) of PC owners had home Internet access, the marginal Internet adopter looked similar to the marginal PC adopter. For households, PC demand had two distinct populations: Those already owning a PC (repeat purchasers), and those that never owned a PC (first-time purchasers). Throughout the 1990s, two distinct Internet adoption patterns correlated with these types of PC demand. Either existing PC adopters converted to the Internet, or households bought PCs and converted to the Internet.

By 2001 to 2002, virtually all existing PC adopters had experience with the Internet at home. Accordingly, the diffusion process changed. There were large differences between existing users and new users in terms of the likelihood of buying a new PC (Prince, 2005). The demand for first-time purchasers was especially relevant since it represented the marginal adopters for PCs, and therefore, strongly resembled the marginal adopters of the Internet.[119]

As the diffusion of the PC moved deeper into mainstream use, the marginal PC adopter became a household with low marginal value for PC quality, high start-up costs, significant price sensitivity, and potential difficulty in determining *when* (not necessarily *if*) to buy. That is why the early Internet adoption experience provided little help for understanding user adoption in later periods. Quite a different set of factors shaped later

[119] In his paper, Prince describes three main determinants of the 'divide' in PC ownership: heterogeneity (in the marginal utility of PC quality and PC holdings), start-up costs, and dynamics. His results indicate that the marginal utility of PC quality is strongly increasing in income and education, and strongly decreasing in age. Further, as prices fall and quality rises over time, the decision about whether to buy a new PC is complicated by the decision of when to buy a new PC. Finally, first time purchasers are more price sensitive than repeat purchasers, and face large start-up costs. See also, Goolsbee and Klenow (2002), who emphasize the role of local network effects in motivating early adoption.

adopter choices than did those for earlier adopters.[120] Similar reasoning partially explains the later appeal of cheaper devices for accessing the Internet, such as netbooks.

While dial-up connection has moved past the frontier stage and approached saturation point in the United States, broadband access approaches the frontier with some frictions preventing uptake. For a few years it was far from ubiquitous, though that is changing as of this writing.[121] As the volume and complexity of traffic on the Internet increases dramatically each year, the value of high-capacity and universal always-on broadband service is constantly increasing. Furthermore, broadband access enables providers to offer a wider range of bundled communications services (e.g., telephone, email, Internet video, etc.) as well as promote more competition between physical infrastructure providers already in place.

In the earliest years of diffusion to households—that is, prior to 2002—the diffusion of broadband Internet access was very much supply-driven in the sense that supply-side issues were the main determinants of Internet availability and, hence, adoption. Cable and telephone firms needed to retrofit existing plants, which constrained availability in many places. In those years, the spread of broadband service was much slower and less evenly distributed than that of dial-up service. Highly populated areas were more profitable due to economies of scale and lower last-mile expenses. As building has removed these constraints, demand-related factors—such as price, bandwidth, and reliability—have played a more significant role in determining the margins between who adopts and who does not.[122]

As of October 2003, 37.2% of Internet users possessed a high-speed connection; the dominant types of broadband access were cable modems and Digital Subscriber Lines (DSL). In addition, broadband penetration has been uneven, as 41.2% of urban and

[120] Such an observation has led to distinct research approaches. For example, Sinai and Waldfogel (2004) investigate whether Internet adoption was motivated by desires to overcome isolation (particularly among minorities) or geographic distance (among those far from retail outlets). For more on urban/rural differences in connectivity see Strover (2001) or Greenstein (2005). Some observers characterize the coming era as not one defined by access to computers, as in the past, but, instead, as one defined by differences in use of computers. Some users will display more sophistication than others, and this will shape differences in returns from investing in computing. See e.g., Hargettai (2003).

[121] Broadband is defined by the FCC as the capability of supporting at least 200 kilobytes per second in at least one direction (supplier and/or consumer), http://www.fcc.gov.

[122] Also see, e.g., Savage and Waldman (2004).

41.6% of central city households with Internet access used broadband compared to a rate of only 25.3% for rural households. Consistent with the supply-side issues, the FCC estimates that high-speed subscribers were present in 97% of the most densely populated zip codes at the end of 2000, whereas they were present in only 45% in the zip codes with the lowest population density (NTIA 2002). By 2006 Internet participation had reached 73% of households, and the supply-side issues began to fade, with only the most low density parts of the country lacking suppliers.

A similar (second) wave of investment occurred in many developed countries over the first decade of the new millennia. Figure 2 shows growth of subscribers per 100 inhabitants in Canada, US, UK, Germany, France, Italy, and Japan, as well as the entire OECD. Though countries differ in the level of broadband use – partly due to household size and other factors, the similarities between them are more apparent. Adoption of broadband grew in all countries.

At this time firms are developing and deploying a wireless delivery channel for some Internet-related services. These options vary in speed, quality, and price. There have been data services from the major cellular carriers (e.g., Verizon, AT&T, and others) since the turn of the millennium, particularly for e-mail delivery to laptops. The most popular mechanism in the recent past was a simple device for delivery of e-mail (e.g., a Blackberry). More complex devices have gained popularity (e.g., iPhones and smart phones), and these have download speeds that begin to approach the low end of wire-line broadband speeds. Technological optimists forecast even faster download speeds from next generation wireless carriers (e.g., WiMax or LTE). There is still considerable uncertainty about how many of these services the market will support, about what price and sales levels will prevail, and, accordingly, what scale of deployment these prices and sales levels will support.

III.7. Co-Invention and Business Processes for Internet Technologies

Some industries are more information intensive than others and, thus, make a more intensive use of new IT developments, such as the Internet, in the production of final goods and services. Heavy computer technology–user industries have historically been banking and finance, utilities, electronic equipment, insurance, motor vehicles, petroleum refining, petroleum pipeline transport, printing and publishing, pulp and paper, railroads, steel, telephone communications and tires (Cortada 1996).

The diffusion of the PC into business did not immediately alter those traditional rankings, but it did introduce computing into some industries that had previously been medium-intensive users, such as warehousing. The constant improvement in the quality of PCs combined with their falling prices along the entire range led to replacement and upgrades of existing systems, as well as to the addition of new uses for the PC.

The growth rates in real investment in computing equipment were extraordinary in the 1990s. Rates of investment in software reached 9.5% growth rates per year for 1990s through 1995, and 14.2% growth rates for 1995 through 2000. Computing equipment growth rates reached, respectively, 13.5% and 7.1% per year per period. Communications equipment growth rates reached 7.2% and 15.5% growth rates. All of these exceeded rates of growth in non-IT capital, which reached 6.8% and 4.9% per year over the same periods.

Business adoption of the Internet was partly responsible for some of the acceleration of investment in the late 1990s; and it came in a variety of forms. By the late 1990s, implementation for minimal applications, such as email, was rather straightforward. It involved a PC, a modem, a contract with an ISP, and some appropriate software. In contrast, investment in the use of the Internet for an application module in a suite of Enterprise Resource Planning software, for example, was anything but routine during the latter half of the 1990s. Such an implementation included technical challenges beyond the Internet's core technologies, such as security, privacy, and dynamic communication between browsers and servers. Usually organizational procedures also changed.[123]

A further motivating factor shaped business adoption: Competitive pressure. That is, there first may be a minimal level of investment necessary just to be in business. Second, there may be investments in the Internet that confer competitive advantage vis-à-vis rivals. Once again, these will vary by locations, industries, and even the strategic positioning of firms (e.g., for example, price leader, high service provider) within those

[123] See Forman and Goldfarb (2006), for a review of studies of Internet investment by business. See Doms (2004) for a review of the acceleration in business investment in PCs.

competitive communities.[124]

Forman, Goldfarb, and Greenstein (2003 a, b) measured national Internet adoption rates for medium and large establishments from all industries. They distinguish between two purposes for adopting, one simple and the other complex. The first purpose, labeled *participation*, relates to activities such as email and Web browsing. This represents minimal use of the Internet for basic communications. The second purpose, labeled *enhancement*, relates to investment in frontier Internet technologies linked to computing facilities. These latter applications are often known as e-commerce, and involve complementary changes to internal business computing processes. The economic costs and benefits of these activities are also distinct; yet, casual analysis in the trade press tends to blur the lines between the two.

They show that adoption of the Internet for purposes of participation is near saturation in most industries. With only a few exceptional, laggard industries, the Internet is everywhere in medium to large business establishments. Their findings for enhancement contrast sharply. There is a strong urban bias towards the adoption of advanced Internet applications. The study concludes, however, that location, per se, does not handicap adoption decisions. Rather, the industries that "lead" in advanced use of the Internet tend to be disproportionately located in urban areas.[125] Related work suggests that small establishments in disadvantaged location may be unable to take advantage of the innovative opportunity due to lack of thick labor markets for technical talent.[126]

III.8. Unending Economic Experimentation

Innovation is experienced by forward looking participants, but understood only in retrospect, and usually only after considerable market experience. It is an exaggeration, but not much of one, to say that before events fully transpire there will be legitimate and passionate debate among participants about which model of value creations most accurately will predict near term events. Only economic experiments can resolve that uncertainty about value.

[124] As Porter (2001) argues, there are two types of competitive motives behind Internet adoption.
[125] See also Forman, Greenstein and Goldfarb (2005).
[126] See Forman, Greenstein and Goldfarb (2008).

As of this writing, the cyclic process of innovative activity continues in computing and the Internet. This can be seen in many places. For example, the dot-com boom has busted and entrepreneurially-led entry has been reborn in a new wave called *web2.0* in sites such as You-Tube, Face Book, and MySpace, which take advantage of social networking among users. The Web is far from done experiencing new waves of entry for new applications, especially in the realm of software applications and other forms of electronic commerce.

In contrast, other layers of the industry have continued to undergo upheaval. Many of the early entrants into the directory business, for example, have lost market share and prominence. Google's innovative approach to searching the Internet and auctioning advertisements next to keywords has largely displaced many existing portals. Even the early leader, Yahoo, has lost some market share in relation to Google.[127]

Another change has begun at the layer of carriers. As users switch from dial-up to broadband Internet access, they also consider switching suppliers. This has led to a large decline in the prominence of AOL as a provider of access to households, and it has led to the ascendancy of the providers of broadband, largely local telephone companies (primarily Verizon, AT&T, and Qwest), and cable firms (primarily Comcast, Cox, Time-Warner, and a few others).[128]

The equipment market has stabilized, leaving Cisco in the dominant position in enterprise computing to serve data communications. Yet, many of Cisco's cousins, firms who grew spectacularly during the 1990s—such as JDS Uniphase, Corning, Lucent, Nortel, and 3Com, did not fare as well. They had to undergo large and painful adjustments in their operations because of the decline in demand (associated with the bursting of the dot-com bubble). Remarkably, the rate of entry of new equipment firms has significantly declined, leaving the incumbent firms more dependent on internal research and development activities than acquisitions (though that began to pick up again as the economy picked up in 2004 and beyond).

Intel also took a more aggressive role in designing PC. Increasingly over the 1990s Intel designed prototypes of these motherboards. By the early 2000 Intel was

[127] These events are discussed more in Haigh (2007).
[128] These events are discussed more in Greenstein (2007).

making some motherboards and encouraging many of its business partners to make similar designs.

In 2003 Intel announced *Centrino*, which marked a departure. It began embedding a Wi–Fi connection in all notebooks that used Intel Microprocessors. To be clear, this *did not* involve redesigning the Intel microprocessor, the component for which Intel is best known. It did, however, involve redesigning the motherboard for desktop PCs and notebooks by adding new parts. This redesign came with one obvious benefit: It eliminated the need for an external card for the notebook, usually supplied by a firm other than Intel, and installed by users (or OEMs—original equipment manufacturers) in an expansion slot. Intel also hoped that its endorsement would increase demand for wireless capabilities within notebooks by, among other things, reducing their weight and size, while offering users simplicity and technical assurances in a standardized function. Intel hoped for additional benefits for users, such as more reliability, less set–up difficulties, and less frequent incompatibility in new settings.

In brief, the Internet is undergoing a second wave of investment. It takes place in the presence of the near saturation of adoption by first-time users. That has been coincident with the presence of more capital deepening in business computing and in deepening operational to support it. These actions are symptomatic that this second wave of investment, unlike the prior one, takes place in the presence of less uncertainty about the source of value.

The new organizational forms for designing new computing also face a series of new tests. Many firms found their business prospects too dependent on Linux to allow it to continue without some structured format, so together they established an firm-sponsored association to continue supporting changes to Linux, employing Linus Torvald and several others in a salaried position. Meanwhile, many new organizations continue to function and support developments, such as Apache, the World Wide Web, and the IETF. Yet, firms no longer naively leave these institutions alone. The standardization organizations find their committees filled with interested participants actively shaping future designs that might affect profitability. These institutions show signs of the stress by slowing down in their decision making, if they reach decisions at all. Perhaps that should also be cause for celebration, since it is a sign that the stakes are high.[129]

III.9. Continuity and Change in Innovative Conduct

The diffusion of a new GPT, the Internet and the World Wide Web, led to a wide variety of changes in computing markets. It was not the technology per se, however, that brought about the most dramatic change. If not technology, what distinguished the latter era from the earlier era?

There were many similarities. In many respects the economic opportunities and challenges resembled those found in prior episodes in computing markets. A new opportunity emerged from efforts at technology push in data networking, and the stretching of the frontier enabled opportunities for value-creation as a GPT diffused. Different firms pursued the economic opportunities they perceived, limited by co-invention costs incurred by both buyers and sellers, and the boundaries of platform competition. Concerns about the emerging platform shaped difference in incumbent and entrepreneurial strategies. In addition, the inherent limits on learning through market experimentation shaped firm understanding about how to create value in this new unexpected commercial opportunity.

Three factors distinguish the Internet era from prior ones. First, the division of technical leadership cut across a wider array of activities than such prior innovative episodes as with the PC and LAN. As a result, a greater variety of market participants reacted to the opportunity. It also made for strange bedfellows. Firms that had little economic relationship to one another prior to the Internet, for example, such as a cable companies and an equipment firm like Cisco, or new firms like CNN and a portal like Yahoo, found themselves making deals and basing their growth projections on the outcomes of those deals. Business assumptions related to the structure of the value chain supporting valuable services and the appropriate innovative conduct for that structure had to be questioned and rethought.

[129] For one interesting account of this slowdown at the IETF, see Simcoe (2007). For an account of the manipulation of hearings at the IEEE, see Mackie-Mason and Netz (2007).

The second distinguishing factor came from the new organizational forms for designing standards in advance of deploying functioning equipment, and, similarly, for altering designs already employed in functioning data networks. The Open Source movement was part of this change, and Linux and Apache, among others, received attention for good reason. Yet, that statement also understates the variety of different organizations for coordinating developments across firms in consensus forums, such as the IETF and IEEE and W3C. These actions changed the boundaries of platform competition, refocusing innovative competition on aspects of products and services other than proprietary features of design. In the short run, the effectiveness of these organizations actually reinforced pre-existing tendencies to specialize in innovative activities as a source of differentiation and strategic advantage. It also raised questions, as yet unanswered, about the durable boundaries between developing proprietary and nonproprietary standards.

The third novel aspect of the commercial Internet involved its breadth of potential applications, leading to a greater breadth of aspiration among participants, and a corresponding change in uncertainty about the source of value from IT across a wide set of participants. The diffusion of the commercial Internet brought about the threat of lasting change in the structure of business, and not all of those changes occurred immediately. The symptoms of this breadth were everywhere in the late 1990s during the dot-com boom and the increasing globalization of production, and still remain in some forms for many firms afterwards. One symptom was that firms had to assign managers to follow developments (that they had previously ignored) in order to make thorough assessments about the direction of change in the source of value. For example, firms in the music distribution business had to follow developments at Internet start-ups, and firms at telephone companies had to understand the implications of the browser wars.

The combination of all three aspects perhaps led to the biggest surprise, widespread exploration by many players in a great many more applications than would have seemed possible or likely only a decade earlier. That is, the managers at firms who never had been big players in IT markets found themselves experimenting with non-proprietary and industry-wide standards-making institutions, facilitating negotiations between firms for key strategic issues. Such institutions co-existed alongside proprietary platforms, sometimes as competitors and sometimes as complements. Cisco, Intel, IBM and many Wi-Fi firms are active participants in these standards forums. Even (previously reluctant) firms such as Microsoft, AT&T, and Verizon have found it useful to participate and fund such activities. Completely new forums, such as the organization built to support Bluetooth, have hundreds of participants.

III.10. Economic Conduct: Open questions

What are the economic determinants of firm conduct? What does the evolution of firm conduct over several decades tell us about those determinants?

Technology push continued to play a role over many decades. Has there been any diminution in the importance market incentives and market-oriented events in computing markets? It appears not. Market forces continue to shape the stretching of the frontier, the severity of change at any point in time, and the identities of leading firms. Co-invention continues to shape the deployment of general purpose technologies, and market forces continue to shape the directions that users and firms pursue.

The experience of the last few decades highlights questions about the causes of variation in the costs of creating value. Why were these costs different for the personal computers, local area networks and the Internet? Why did they differ over time? What economic mechanisms link such costs link to the size of the economic benefits from deploying new technology? These are some of the major unanswered questions highlighted by these events.

There also has been no decline in the economic role of computing platforms. Platforms shape economic behavior today as much as in any other time in the past. In part this happened because market participants have learned about the importance of platform strategies for creating value for firms, such as Intel and Microsoft. In part, changes in the unification or division of technical leadership within and between platforms has played a role in shifting value from one set of firms to another, as from one Yahoo! to Google. How does the presence of platforms color the economic incentives of leading incumbent firms and new entrants to innovate? What mechanisms shape those incentives, and do these push firms to undertake innovation that reinforces or alters market structure? These are unanswered questions.

The role of market-based learning activity also has not diminished, but the economic behavior of the last few decades raises questions about the determinants of its prevalence and importance over time. How does behavior aimed at reducing uncertainty – such as the prevalence of platforms and localization of innovative conduct in a small set of locations – alter learning activities? Will other factors, such as the globalization of production and use, make substantial differences to related conduct? Similar questions apply to the unprecedented dispersion of uncoordinated innovative conduct that characterized the early diffusion of the Internet. Was that behavior merely an artifact of early diffusion or does it portends something like a permanent change in the norms for competitive behavior, resulting from coordinating participants with disparate commercial interests, such as found in open source platforms?

Another set of unanswered questions relate to the wide breadth of economic activity touched by computing. Many economic participants have built a network with a high degree of technical interrelatedness. Will the general gains to all parties from bringing routines into business activities overwhelm the discretion of innovative actors, limiting their innovative conduct? For the time being there appears to be no cessation in the never-ending nature of the investment in economic experiments, so that seems to imply that any breakthrough in widely used IT could have large economic effects on the entire economy.

Experimentation did not end with just one episode or the end of the dot-com bubble. Indeed, as of this writing, many questions remain open about the value of different aspects of IT in the long run, and firms continue to explore approaches to creating value. What is particularly notable is the lack of cessation of technically-oriented entrepreneurialism. In the not so distant pass some participants had expected to undertake economic experiments all along, while other participants had previously thought they did not live in an entrepreneurial technology market. Instead, different expectations have melted away. Virtually all participants in these markets expect continual entrepreneurially led change, as well as its twin, the absence of economic tranquility.

References

Abbate, Janet, 1999, *Inventing the Internet*, MIT Press; Cambridge, MA.

Acemoglu, Daron, Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizzio Zilibotti, 2007, "Technology, Information and the Decentralization of the Firm," *Quarterly Journal of Economics.*

Agarwal, Ajay, and Avi Goldfarb, 2006, "Restructuring Research: Communications Costs and the Democratization of University Innovation," Working paper, University of Toronto..

Aizcorbe, Ana, 2006. "Why did Seminconduction Price indexes Fall So Fast in the 1990s?" *Economic Inquiry*, 44(3), pp. 485-496, July.

Aizcorbe, Ana, Flamm, Ken, and Anjum Khursid, 2007, "The role of Semiconductor Inputs in IT Hardware Price Decline: Computers vs. Communications," in (eds) Ernst R. Berndt and Charles M. Hulten, *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, University of Chicago Press.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella, 2001, *Markets for Technology: The Economics of Innovation and Corporate Strategy*, MIT Press.

Arora, Ashish and Farasat Bokhari, 2007, "Open versus Closed firms and the Dynamics of Industry Evolution," *Journal of Industrial Economics*

Arora, Ashish, Jon Caulkins and Rahul Telang, 2006, "Sell first, patch later: Impact of patching on software quality," *Management Science*, 52:3, pp. 465-471.

Arora, Ashish, and Alfonso Gambardella, 2005, *From Underdogs to Tigers: The Rise and Growth of the Software Industry in Brazil, China, India, Ireland and Israel.* Oxford University Press.

Arora, Ashish, and Chris Forman, 2006, "How Local are IT Outsourcing Markets: Proximity and Software Programming," Available at http://ssrn.com/abstract=953254

Aspray, William, and Martin Campbell-Kelly, 1996, *Computer: A History of the Information Machine*, Basic Books; New York.

Astebro, Thomas. 2002. "Noncapital investment costs and the adoption of CAD and CNC in U.S. metalworking industries" *RAND Journal of Economics* 33(4): 672-688.

Astebro, Thomas, 2004, "Sunk Costs and the Depth and Probability of Technology Adoption," *Journal of Industrial Economics*, LII(3), 2004, 381-399

Athey, S. and Stern S., 2002, "The Impact of Information Technology and Job Design on Emergency Health Care Outcomes," *RAND Journal of Economics,*

Atrostic, Barbara K., & Sang V. Nguyen, 2005. "IT and Productivity in US Manufacturing: Do Computer Networks Matter?" *Economic Inquiry,* vol. 43(3), pages 493-506, July.

Attewell, Paul. 1992. "Technology Diffusion and Organizational Learning: The Case of Business Computing." *Organizational Science* 3 (February): 1-19.

Augereau, Angelique, Shane Greenstein, and Marc Rysman, 2006, "Coordination versus Differentiation in a Standards War: 56K modems." *Rand Journal of Economics.* 34 (4), pp. 889-911.

Autor, David, Lawrence F. Katz, and Melissa S. Kearney, 2003, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118 (4), pp. 1279-1334.

Baker, George, and Tom Hubbard, 2003, "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information," *American Economic Research*, June, pp. 551-572.

Baker, George, and Tom Hubbard, 2004, "Contractibility and Asset Ownership: On-Board Computers and Governance in US Trucking", *Quarterly Journal of Economics*, November, pp. 1443-1480.

Baker, Jonathan, 2007, "Beyond Schumpeter vs. Arrow: How Antitrust Fosters Innovation," Paper presented at NBER Summer Institute, Innovation Policy and the Economy, Accessed at http://nber15.nber.org/~confer/2007/si2007/PRIPE/baker.pdf.

Bakos, Yannis, and Brynjolfsson, Erik (2000) Bundling and Competition on the Internet: Aggregation Strategies for Information Goods, *Marketing science,* Vol. 19, No. 1 pp. 63-82.

Bakos, Yannis, Brynjolfsson, Erik and Lichtman, Douglas (1999) Shared Information Goods, *Journal of Law and Economics*, Vol. 42, No. 117.

Baldwin, Carliss A., and Kim B. Clark, 1997, "Sun Wars: Competition within a Modular Cluster," in *Competing in the Age of Digital Convergence*, (Ed) D. B. Yoffie, Boston, MA: Harvard Business School Press.

Baldwin, Carliss A., and Kim B. Clark, 2006, "Architectural Innovation and Dynamic Competition: The Smaller Footprint Strategy," Working Paper. Harvard Business School.

Bank, David. 2001. *Breaking Windows: How Bill Gates Fumbled the Future of Microsoft.* London: The Free Press.

Beaudry, Paul, Mark Doms, Ethan Lewis, 2006, "Endogenous Skill Bias in Technology Adoption: City-Level Evidence from the IT Revolution," *Federal Reserve Bank of San Francisco, Working paper 06-24.*

Berndt, Ernst R., Zvi Griliches and Neal Rappaport, 1995, "Econometric Estimates of Price Indexes for Personal Computers in the 1990's," *Journal of Econometrics*, Vol. 68, No. 1, July, pp. 243-268.

Berners-Lee, Tim, and Mark Fischetti, 1999, *Weaving the Web, The Original Design andUltimate Destiny of the World Wide Web*, Harper Collins, New York.

Bloom, Nick, and John Van Reenen (2007), "Measuring and explaining management practices across firms and countries," *Quarterly journal of economics*, 122 (4). pp. 1351-1408.

Bloom, Nick, Rafaella Sadun, and John Van Reenen, 2007, "Americans do I.T. Better: U.S. Multinationals and the Productivity Miracle," CEP Discussion Paper 788, April.

Blumenthal, Marjory S., and David D. Clark, 2001, "Rethinking the Design of the Internet: The End-to-End Arguments vs. The Brave New World." In (eds.) Benjamin Compaine and Shane Greenstein, *Communications Policy in Transition: The Internet and Beyond.* MIT Press: Cambridge, MA. Pp. 91-139.

Joseph Bower and Clayton Christensen, 1997, "Disruptive Technologies: Catching the Wave" in (Ed) John Seely Brown, Seeing Differently: Insights on Innovation, *Harvard Business Review Book.*

Bresnahan, Timothy, 1999, "The Changing Structure of Innovation in Computing" in (ed) Jeffrey A Eisenach and Thomas M Lenard, *Competition, Convergence and the Microsoft Monopoly: Antitrust in the Digital Marketplace*, Kluwer Academic Publishers, Boston.

Bresnahan, Timothy, 2004, "The Economics of the Microsoft Case," Working paper, accessed at http://www.stanford.edu/~tbres/research.htm.

Bresnahan, Timothy, Erik Brynjolfsson, and Lorin Hitt, 2002, Information Technology, Work Organization, and the Demand for Skilled Labor: Firm-Level Evidence. Quarterly *Journal of Economics* 117 (February): 339-376.

Brensnahan, Timothy, and Shane Greenstein, 1997, "Technical Progress and Co-Invention in Computing and in the Use of Computers." *Brookings Papers on Economics Activity: Microeconomics*, Pp. 1-78.

Bresnahan, Timothy and Shane Greenstein, 1999, "Technological Competition and the Structure of the Computer Industry," *Journal of Industrial Economics*, March, pp 1-40.

Bresnahan, Timothy, Shane Greenstein, and Rebecca Henderson, 2006, “Making Waves: The Interplay between Market Incentives and Organizational Capabilities in the Evolution of Industries,” Paper presented at Harvard Business School Strategy Conference, October.

Brensnahan, Timothy, and Franco Malerba, 1999, "Industrial Dynamics and the Evolution of Firm’s and Nations’ Competitive Capabilities in the World Computer Industry," *Sources of Industrial Leadership*, D. Mowery and R. Nelson, Eds, Cambridge University Press, 1999.

Bresnahan, Timothy, and Manuel Trajtenberg, 1995. “General Purpose Technologies: ‘Engines of growth’?” *Journal of Econometrics*, 65 (1): 83-108.

Bresnahan, Timothy, Scott Stern, and Manuel Trajtenberg, 1997, “Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s The *RAND Journal of Economics* 28(0), pp. S17-S44.

Bresnahan, Timothy, and Pai-Ling Yin, 2007, “Standard Setting in Markets: The Browser Wars,” in (eds) Shane Greenstein and Victor Stango, *Standards and Public Policy,* Cambridge University Press; Cambridge, UK. Pp. 18-59.

Brynjolfsson, Erik and Hitt, Lorin, 2000, “Beyond Computation: Information Technology, Organizational Transformation and Business Performance,” *Journal of Economic Perspectives*, Vol. 14, No. 4, pp. 23-48.

Brynjolfsson, Erik, Hitt, Lorin M. and Shinkyu Yang (2002) "Intangible Assets: Computers and Organizational Capital," *Brookings Papers on Economic Activity*: Macroeconomics (1): 137-199.

Card, David and John DiNardo, 2002, “Skill-Biased Technical Change and Rising Wage Inequality: Some Problems and Puzzles,” *Journal of Labor Economics*, 20 (4), pp. 733- 83.

Caroli, Eve, and John Van Reenen, 2001, “Sill Biased Organizational Change? Evidence from the British and French Establishments,” *Quarterly Journal of Economics,* CXVI, No. 4, 1449-1492.

Carr, Doug, 1999, *IBM Redux: Lou Gerstner and the Business Turnaround of the Decade.* Harper Collins; New York.

Chapman, Merrill R., 2006, *In Search of Stupidity: Over Twenty Years of High*

Tech Marketing Disasters Springer-Verlag: New York.

Christensen, Clay, 1997, *Innovators Dillemma.* Harvard Business School Press; Boston, MA.

Cockburn, Iain, and Megan MacGarvie, 2006, "Entry, Exit and Patenting in the Software Industry," NBER Working Paper No. W12563.

Compaine, Benjamin, and Shane Greenstein, 2001, *Communications Policy in Transition: The Internet and Beyond.* MIT Press: Cambridge, MA.

Cortada, James W., 1996. *Information Technology as Business History: Issues in the History and Management of Computers.* Westport, CT: Greenwood Press.

Cortada, James W., 2003, *The Digital Hand: How Computers Changed the Work of American Manufacturing, Transportation and Retailing Industries*, Cary, NC: Oxford University Press.

Cranor, Lorrie, and Shane Greenstein, 2002, *Communications Policy and Information Technology: Promises, Problems, Prospects*, MIT Press. 2002

Cranor, Lorrie and Steve Wildman, 2003, *Rethinking Rights and Regulations: Institutional Responses to New Communications Technologies*, Cambridge, MIT Press.

Cringley, Robert X., 1992, *Accidental Empires*, Harper Collins, New York.

Cusumano, Michael, and Richard Selby, 1995, *Microsoft Secrets, How the World's Most Powerful Software Company Creates Technology, Shapes Markets and Manages People,* Simon and Schuster Inc; New York.

Cusumano, Michael, and David Yoffie, 2000, *Competing on Internet Time: Lessons from Netscape and its Battle with Microsoft*, Free Press.

Dalle, Jean-Michel, Paul A. David, Rishab Aiyer Ghosh, and Frank Wolak, 2004, "Free & Open Source Software Developers and 'the Economy of Regard': Participation and Code-Signing in the Modules of the Linux Kernel," Working paper, SIEPR, Stanford University, Open Source Software Project, Accessed on June, 2007, at http://siepr.stanford.edu/programs/OpenSoftware_David/NSFOSF_Publications.html

David, Paul, 1990, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review*, 80, 2.

David, Paul, 2001, "The Evolving Accidental Information Super Highway," *Oxford Review of Economic Policy*, 17 (2), pp. 159-187.

Dedrick, Jason, and Joel West, 2001, "Open Source Standardization: the Rise of Linux in the Network Era," *Knowledge, Technology and Policy*, V 14 (2), pp 88-112.

Dedrick, Jason, and Kenneth L. Kraemer, 2005, "The Impacts of IT on Firm and Industry Structure: the Personal Computer Industry," *California Management Review,* 47(3), pp 122-142.

De Figueiredo, John M., Margaret K. Kyle, 2006, "Surviving the gales of creative destruction: the determinants of product turnover," *Strategic Management Journal*, 27 (3), pp. 241-264.

Doms, Mark 2004, "The Boom and Bust in Information Technology Investment," *Federal Reserve Bank of San Francisco Economic Review,* pp 20 – 34.

Downes, Tom, and Shane Greenstein, 2002 "Universal Access and Local Internet Markets in the U.S." *Research Policy,* 31: 1035-52.

----, 2007, "Understanding Why Universal Service Obligations May Be Unnecessary: The Private Development of Local Internet Access Markets" *Journal of Urban Economics.*

Draca, Mirko, Raffaella Sadun, John Van Reenen, 2007, "ICT and Productivity," in (eds) Robin Mansell, Chrisanthi Avgerou, Danny Quah, Roger Silverstone, *Oxford Handbook of ICTs,* Oxford University Press; Oxford, UK.

Dranove, David, and Neil Gandal, 2003, "The DVD vs. DIVX Standard War. Empirical Evidence of Network Effects and Preannouncement Effects," *Journal of Economics and Management Strategy,* 12 (3), pp. pp. 363-386.

Fallick, Bruce, Charles Fleischman, Charles, and James A. Rebitzer, 2006. "Job-Hopping in Silicon Valley: Some Evidence Concerning the Micro-foundation of a High-Technology Cluster," *Review of Economics and Statistics*, 88(3): 472-481.

Farrell, Joseph, and Garth Saloner,1985. "Standardization, Compatibility, and Innovation, *Rand Journal of Economics,* Vol. 16(1), pages 70-83, Spring.

Ferguson, Charles, 1999, *High Stakes, No Prisoners: A winners Tale of Greed and Glory during the Internet Wars.*, Times Books.

Fichman, Robert and Chris Kemerer. 1997. "The Assimilation of Software Process Innovations: An Organizational Learning Perspective." *Management Science* 43(10): 1345-1363.

Fisher, Franklin, John J. McGowan, and J.E. Greenwood, 1983, *Folded Spindled and Mutilated: Economic Analysis and U.S. vs. IBM,* MIT Press.

Fisher, Franklin M., J. W. McKie, and R. B. Mancke, 1983, *IBM and the US Data Processing Industry, An Economic History*, Praeger Publishers, New York.

Fisher, Franklin M., 2000, "The IBM and Microsoft Cases: What's the Difference?" *American Economic Review*, 90 (May 2000): 180-183

Flamm, Kenneth, 2003, Chapter 2, "The New Economy in Historical Perspective: Evolution of Digital Electronics Technology," in D.C. Jones, Ed., *New Economy Handbook* (San Diego: Academic Press/Elsevier).

Flamm, Kenneth, 1987, *Creating the Computer: Government, Industry and High Technology*, The Brookings Institution. Pp. 203-234.

Forman, Chris, 2005. "The Corporate Digital Divide: Determinants of Internet Adoption," *Management Science*, 51(4): 641-54.

Forman, Chris, and Avi Goldfarb, 2006, "Diffusion of Information and Communications Technology to Business," in (ed) Terry Hendershott, *Economics and Information Systems, Volume 1,* Elsevier.

Forman, Chris, Avi Goldfarb, and Shane Greenstein, 2003a, "Which Industries Use the Internet?" in (ed) Michael Baye *Organizing the new Industrial Economy*, Elsevier, pp 47-72.

–––, 2003b, "The Geographic Dispersion of Commercial Internet Use." in (ed) Lorrie Cranor and Steve Wildman, *Rethinking Rights and Regulations: Institutional Responses to New Communications Technologies*, Cambridge, MIT Press, 113-45.

–––, 2005. "How did Location Affect Adoption of the Internet by Commercial Establishments? Urban Density versus Global Village," *Journal of Urban Economics*, 58(3): 389-420.

–––, 2008. "Understanding Inputs into Innovation: Do Cities Substitute for Internal Firm Resources?" *Journal of Economics and Management Strategy*. Pp 295-316.

Fosfuri, Andrea, Marco Giarratana, and Alessandra Luzzi, 2005, "Firm Assets and Investments in Open Source Software Products," Druid Working Paper No 05-10. Copenhagen Business School.

Franco, April, and Darren Filson, 2006, "Spin-outs: Knowledge Diffusion Through Employee Mobility," *Rand Journal of Economics*, 37(4), pp. 841-60.

Frieberger, P., and M. Swaine, 1984, *Fire in the Valley: The Making of the Personal Computer,* Osborne/McGraw Hill, Berkeley, CA.

Gandal, Neil, Shane Greenstein, and David Salant, 1999, “Adoption and Orphans in the Early Computer Market: Chicken and Eggs, Hardware and Software,” *Journal of Industrial Economics*, pp 87-105.

Gandal, N., M. Kende, and R. Rob, 2000, "The Dynamic of Technological Adoption in Hardware/Software Systems: The Case of Compact Disc Players," *Rand Journal of Economics*, 31, 43–61.

Gans, Joshua, David Hsu, and Scott Stern, 2002, “When Does Start-up Innovation Spur the Gale of Creative Destruction?” *Rand Journal of Economics*, 33(4), pp. 571-586.

Gans, Joshua, and Scott Stern, 2003, “The Product Market and the Market for Ideas: Commercialization Strategies for Technology Entrepreneurs,” *Research Policy*, 32(2), pp. 333-350.

Gates, Bill, 1995, *The Internet Tidal Wave*, Government Exhibit 20, United States versus Microsoft, accessed at http://www.usdoj.gov/atr/cases/ms_exhibits.htm.

Gawer, Annabelle, and Michael Cusumano, 2002, *Platform Leadership: How Intel, Microsoft and Cisco Drive Innovation,* Harvard Business School Press; Boston, MA.

Gawer, Annabelle, and Rebecca Henderson, 2007, “Platform Owner Entry and Innovation in Complementary Markets: Evidence from Intel,” *Journal of Economics and Management Strategy*, Volume 16 (1).

Gerstner, Louis V., 2002, *Who says Elephants Can’t Dance? Leading a Great Enterprise Through Dramatic Change*, Harper Collins; New York.

Gilbert, Richard J., 1998, “Networks, Standards, and the Use of Market Dominance: Microsoft (1995),” in (eds) John Kwoka and Larry White, *The Antitrust Revolution: The Role of Economics*, 3rd edition, Oxford: Oxford University Press.

Gilbert, Richard J., 2006, “Looking for Mister Schumpeter: Where are we in the Competition-Innovation Debate?” in (eds) Adam Jaffe, Josh Lerner, and Scott Stern, *Innovation Policy and the Economy,* Volume 8, pp 159.

Gilles, James, and Robert Cailliau, 2000, *How the Web was Born,* Oxford University Press.

Goldfarb, Avi, 2004. “Concentration in Advertising-Supported Online Markets: An Empirical Approach.” *Economics of Innovation and New Technology*, 13(6), 581-594.

Goldfarb, Brent D., Kirsch, David and Pfarrer, Michael D., 2005, "Searching for Ghosts: Business Survival, Unmeasured Entrepreneurial Activity and Private Equity Investment in the Dot-Com Era". Robert H. Smith School Research Paper No. RHS 06-027 Available at SSRN: http://ssrn.com/abstract=825687

Goolsbeee, Austin, and Peter Klenow, 2002, “Evidence on Learning and Network Externalities in the Diffusion of Home Computers,” *Journal of Law and Economics*, October, Vol XLV (2, part 1), 317-344.

Gordon, R. J., 2000, “Does the New Economy Measure up to the Great Inventions of the Past?” *Journal of Economic Perspectives*, 14, 4.

Greenstein, Shane, 2000, “Building and Delivering the Virtual World,” *Journal of Industrial Economics.* December. Pp 391-411.

Greenstein, Shane, 2001, “Commercialization of the Internet: The Interaction of Public Policy and Private Action,” in (eds) Adam Jaffe, Josh Lerner, and Scott Stern, *Innovation Policy and the Economy.*

Greenstein, Shane, 2005, “The Economic Geography of Internet Infrastructure in the United States,” in (eds) Martin Cave, Sumit Majumdar, and Ingo Vogelsang, *Handbook of Telecommunication Economics, Volume II.* Elsevier.

Greenstein, Shane, 2006, “Wikipedia in the Spotlight.” Kellogg School of Management, Northwestern University, Teaching Case, accessed at http://www.kellogg.northwestern.edu/faculty/greenstein/images/cases.htm.

Greenstein, Shane, 2007a “"Economic Experiments and Neutrality in Internet Access Markets," in (eds) Adam Jaffe, Josh Lerner and Scott Stern, *Innovation, Policy and the Economy, Volume 8,* MIT Press.

Greenstein, Shane, 2007b, “The Evolution of Market Structure for Internet Access in the United States.” in (eds) William Aspray and Paul Ceruzzi, *The Commercialization of the Internet and its Impact on American Business*, MIT Press.

Greenstein, Shane, and Ryan McDevitt. 2009. The Broadband Bonus: Accounting for Broadband Internet’s Impact on U.S. GDP. NBER Working paper 14758. http://www.nber.org/papers/w14758.

Greenstein, Shane, and Jeff Prince, 2007, “The Diffusion of the Internet and the Geography of the Digital Divide,” in (Eds. Robin Mansell, Danny Quah, and Roger Silverstone), *Oxford Handbook on ICTs*, Oxford University Press

Greenstein and Wade, 1998, "The Product Life Cycle in the Commercial Mainframe Computer Market, 1968-1983," *Rand Journal of Economics*.

Grove, Andrew, 1996, *Only the Paranoid Survive: How to Exploit the Crisis Points That Challenge Every Company*, Doubleday; New York.

Haigh, Thomas, 2007, "Building the Web's Missing Links: Portals and Search

Engines" in *The Internet and American Business*, in (eds) William Aspray and Paul Ceruzzi, *The Internet and American Business,* MIT Press.

Hargettai, E., 2003, 'The Digital Divide and What To Do About It', in D. C. Jones (ed.) *The New Economy.* San Diego, CA: Academic Press, 822-38.

Harrison, Rupert, Rachel Griffith, and John Van Reenen, 2006, "How special is the relationship: Using the impact of the US R&D spillvers on British firms as a test of technology sourcing," *American Economic Review* (December) 96(5), 1859-1875

Henderson, Rebecca, 2000, Declaration of Rebecca Henderson, United States v. Micros oft Remedies Papers, http://www.usdoj.gov/atr/cases/f4600/4644.htm

Henry, David and Donald Dalton, 2002, "Information technology industries in the new economy," in Digital Economy 2002, Chapter III. Washington DC: Department of Commerce, http://www.esa.doc.gov/reports.cfm.

Hoetker, Glenn, 2006, "Do Modular Products Lead to Modular Organizations?" *Strategic Management Journal,* 27: 501: 518.

Hubbard, Tom, 2000, "The Demand for Monitoring Technologies: The Case of Trucking," *Quarterly Journal of Economics,* May, pp 533-560.

Hubbard, Tom, 2003, "Information, Decisions, and Productivity: On-Board Computers and Capacity Utilization in Trucking," in *American Economic Review,* September.

Hussain, Farooq, 2003, "Historic Role of the Commercial Internet eXchange Router and its Impact on the Development of the Internet eXchange Points [IXCs]," http://www.farooqhussain.org/projects/cixrouter19912001/index_html,accessed at July, 2007.

Jorgenson, Dale W., 2001, "Information Technology and the U.S. Economy," *American Economic Review*, 91, 1, March, pp. 1 – 32.

Jorgenson, Dale W., Mun S. Ho, and Kevin Stiroh, 2005, *Productivity, Volume 3: Information Technology and the American Growth Resurgence,* MIT Press.

Jorgenson, Dale, W., and Charles W. Wessner, Editors, 2005, *Deconstructing the Computer, Report of a Symposium,* National Academies Press: Washington, D.C.

Kahin, Brian, and James Keller, 1995, *Public Access to the Internet*, MIT Press; Cambridge, MA.

Kahin, Brian, and James Keller, 1997, *Coordinating the Internet,* MIT Press; Cambridge, MA.

Kahin, Brian and Bruce McConnell, 1997, "Towards a Public Metanetwork; Interconnection, Leveraging and Privatization of Government-Funded Networks in the United States," in (eds) Eli Noam and Aine Nishuilleabhain, *Private Networks Public Objectives,* Elsevier, Amsterdam, the Netherlands. Pp. 307 – 321.

Kahn, Robert, 1995, 'The Role of Government in the Evolution of the Internet,' in National Academy of Engineering (ed.) *Revolution in the U.S. Information Infrastructure*. Washington, DC: National Academy Press, 13-24.

Katz B., and Phillips, A., 1982, "The Computer Industry," in Nelson R. R. (ed), *Government and Technical Progress: A Cross-Industry Analysis* (Pergamon Press, New York), pp 162-232.

Katz, Michael, and Carl Shapiro, 1985, "Network Externalities, Competition and Compability," *American Economic Review,* 75 (3), pp. 424-440.

Michael L. Katz and Carl Shapiro, 1986. "Technology adoption in the presence of network externalities." *Journal of Political Economy,* 94(4):822–84.

Michael L. Katz and Carl Shapiro, 2000. "Systems competition and network effects." *Journal of Economic Perspectives,* 8(2):93–115, 1994

Kende, Michael, 2000, "The Digital Handshake: Connecting Internet Backbones," Working Paper No. 32, Federal Communications Commission, Office of Planning and Policy, Washington, DC.

Kenney, Martin, 2000, *Understanding Silicon Valley: The Anatomy of an Entrepreneurial Region,* Stanford University Press, Stanford, Ca.

Kenney, Martin, and Urs Von Burg, 1999, "Technology, Entrepreneurship, and Path-Dependence: Industrial Clustering in Silicon Valley and Route 128," *Industrial and Corporate Change,* 8, pp. 67-103.

Kesan, Jay P., and Rajiv C. Shah, 2001, "Fool us Once, Shame on you – Fool us Twice, Shame on us: What we Can Learn from the Privatizations of the Internet Backbone Network and the Domain Name System," *Washington University Law Quarterly,* 79, pp 89 – 220.

Khanna, T., 1995, "Racing Behavior: Technological Evolution in the High-end Computer Industry," *Research Policy* 24, no. 6.

Kreuger, Alan, 1993, "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989.," *Quarterly Journal of Economics,* 108 (1), pp. 33-60.

Langlois, Richard, and David Mowery, 1996, “Spinning off and Spinning on (?): The Federal Government Role in the Development of the U.S. Computer Software Industry,” *Research Policy,*

Langlois R., and P. L. Robertson, 1992, "Networks and Innovation in a Modular System: Lessons from the Microcomputer and Stereo Component Industries," *Research Policy*, 21, PP 297-313.

Langlois R., and P. L. Robertson, 1995, *Firms, Markets and Economic Change: A Dynamic Theory of Business Institutions,* Routledge.

Lee, Chong-Moon, William F Miller, Marguerite Gong Hancock, and Henry S. Rowen, 2000, *The Silicon Valley Edge: A Habitat For Innovation and Entrepreneurship,* Stanford University Press, Stanford, CA.

Lerner, J., 1997, "An Empirical Analysis of a Technology Race," *Rand Journal of Economics*, 28: 228-247.

Lerner, Josh, and Jean Tirole, 2002, “Some Simple Economics of Open Source Software,” *Journal of Industrial Economics,* 50 (2), pp. 197-234.

Lucking-Reily, David, and Daniel Spulber, 2001, “Business to Business Electronic Commerce,” *Journal of Economic Perspectives,* Volume 15 (1), pp. 55-68.

Mackie-Mason, Jeff, and Janet Netz, 2007, “Manipulating Interface Standards as an Anticompetitive Strategy,” in (eds) Shane Greenstein and Victor Stango, *Standards and Public Policy,* Cambridge Press; Cambridge, MA. Pp. 231-259.

Mankiw, Gregory, Kristin Forbes, and Harvey Rosen, 2005, “Chapter 6. Innovation and the Information Economy,” in *Economic Report of the President*, US Government Printing Office, Washington.

Mansell, Robin, and W. Edward Steinmueller, 2000, *Mobilizing the Information Society Strategies for Growth and Opportunity,* Oxford University Press.

McKinsey Global Institute, 2001, *U. S. Productivity Growth, 1995-2000: Understanding the Contribution of Information Technology Relative to Other Factors,* McKinsey and Co., Washington D.C.

Meeker, Mary, and Chris DuPuy, 1996, *The Internet Report,* Harper Business, New York.

Messerschmitt and Szyperski, 2003, *Software Ecosystems: Understanding an Indespensible Technology and Industry,* MIT Press.

Mockus, Audris, Roy T. Fielding, and James D. Herbsleb, 2005, "Two Case Studies of Open Source Software Development: Apache and Mozilla," in (ed) Joseph Feller, *Perspectives on Free and Open Source Software,* MIT Press; Cambridge.

Mowery, David, 1996, *The International Computer Software Industry, A Comparative Study of Industry Evolution and Structure,* Oxford University Press.

Moore, Gordon, 1965, "Cramming More Components onto Integrated Circuits," Electronics, V38 (8), April, pp 114-17.

Mowery, David C. and Simcoe, Tim S., 2002a, 'The Origins and Evolution of the Internet,' in R. Nelson, B. Steil, and D. Victor (eds) *Technological Innovation and Economic Performance.* Princeton, NJ: Princeton University Press, 229-64.

--- 2002b, "Is the Internet a U.S. Invention? An Economic and Technological History of Computer Networking," *Research Policy.*

Mueller, Milton L., 2002, *Ruling the Root: Internet Governance and the Taming of Cyberspace,* MIT Press; Cambridge, MA.

National Research Council, 2003, *Innovation in Information Technology,* National Academy Press; Washington, D.C.

NTIA (National Telecommuncations and Information Administration), 2004. 'A Nation Online: Entering the Broadband Age', http://www.ntia.doc.gov/reports.html.

—, 2002, 'A Nation Online: How Americans Are Expanding Their Use Of The Internet', http://www.ntia.doc.gov/reports.html, assessed 17 Mar. 06.

—, 1995, 'Falling Through the Net: A Survey of the 'Have Nots' in Rural and Urban America', http://www.ntia.doc.gov/reports.html, assessed 17 Mar. 06.

Noam, Eli, 2001, *Interconnecting the Network of Networks.* Cambridge, MA: MIT Press.

Nuechterlein, J. E. and Weiser, P. J., 2005, *Digital Crossroads: American Telecommunications Policy in the Internet Age*, Cambridge, MA: MIT Press.

Ohashi, Hiroshi, 2003, "The Role of Network Effects in the US VCR Market, 1978-1986," *Journal of Economics & Management Strategy*, 12 (4), pp. 447–494.

Ou, George, 2008, A Policy Maker's Guide to Network Management. *The Information Technology and Innovation Foundation.* http://www.itif.org/index.php?id=205.

Oxman, J., 1999, 'The FCC and the Unregulation of the Internet,' Working paper

31, Federal Communications Commission, Office of Planning and Policy, Washington, DC.

Partridge, Craig. 2008. The technical development of Internet e-mail. *IEEE Annals of the History of the Computing*, April, 3-29.

Paulsen, Ed, 2001, *Inside Cisco: The Real Story of Sustained M&A Growth*, John Wiley and Sons; New York.

Porter, Michael, 2001, 'Strategy and the Internet'. *Harvard Business Review*. 79(3): 62-78.

Prince, Jeff, 2005, 'Measuring the Digital Divide: Structural Estimation of the Demand for Personal Computers', mimeo, Cornell University.

Prusa, Thomas J., James A. Schmitz, Jr., 1994, "Can Companies Maintain their Initial Innovative Thrust? A Study of the PC Software Industry," *The Review of Economics and Statistics, Vol*. 76, No. 3., pp. 523-540.

Quarterman, John S., 1989, *Matrix Computer Networks and Conferences*, Digital Press.

Rohlfs, Jeffrey, 2001, *Bandwagon Effects in High Technology Industries,* MIT Press; Cambridge, MA.

Rosenberg, Nathan, 1976, *Perspectives on Technology,* Cambridge University Press; Cambridge, UK.

Rosenberg, Nathan, 1994, "Economic Experiments," in *Exploring the Black Box,* Cambridge University Press, Cambridge.

Rosenberg, Nathan, 1995, "Uncertainty and Technology Change," in (eds.), Landau, Taylor and Wright, *The Mosaic of Economic Growth*, Stanford University Press, Stanford, CA. 1996. Pp. 334-356.

Russell, Andrew L., 2006, "Rough Consensus and Running Code and the Internet-OSI Standards War," *IEEE Annals of the History of Computing*, July-September, pp. 48-61.

Sang Nguyen & B.K. Atrostic, 2006. "How Business Use Information Technology: Insights for Measuring Technology and Productivity," Working paper 06-15, Center for Economic Studies, U.S. Census Bureau.

Savage, Scott, and Donald Waldman, 2004, "United States Demand for Internet Access," *Review of Network Economics*, 3(3), pp. 228-247.

Saxenian, Annalee, 1994, *Regional Advantage: Culture and Competition in Silicon Valley and Route 128,* Cambride, Mass; Harvard University Press, 1994.

Schmalensee, Richard, 2000, "Antitrust Issues in Schumpeterian Industries." *American Economic Review*, 90 (May), pp. 192-196.

Shapiro, Carl, and Hal R. Varian, 1998, *Information Rules.* Harvard Business School Press.

Shiller, Robert J., 2000, *Irrational Exuberance*, Princeton University Press; New York.

Simcoe, Timothy, 2007, "Delay and De jure Standardization: Exploring the Slow Down in Internet Standards Development," in (eds) Shane Greenstein and Victor Stango, *Standards and Public Policy*, Cambridge Press; Cambridge, MA. Pp. 260-297.

Sinai, T. and Waldfogel, J., 2004. 'Geography and the Internet: Is the Internet a Substitute or a Complement for Cities?' *Journal of Urban Economics* 56(1): 1-24.

Sink, Eric, 2007, "Memoirs from the browser wars." Published at http://biztech.ericsink.com/Browser_Wars.html, accessed March, 2007.

Smith, Michael.D., Bailey, Joseph, and Brynjolfsson, Erik (2000) Understanding Digital Markets: Review and Assessment, MIT Press, pp. 276-288. in Brynjolfsson, Erik and Kahin, Brian, *Understanding the Digital Economy*, MIT Press.

Stavins, Joanna, 1995, Model Entry and Exit in a Differentiated-Product Industry: The Personal Computer Market , *The Review of Economics and Statistics*, Vol. 77, No. 4. pp. 571-584.

Stern, Scott, 2005, "Economic Experiments: The Role of Entrepreneurship in Economic Prosperity," in *Understanding Entrepreneurship: A Research and Policy Report.* http://research.kauffman.org/cwp/ShowProperty/web/CacheRepository/Documents/Research__Policy_Singles.pdf, Kaufman Foundation.

Stranger, Greg, and Shane Greenstein (2007), "Pricing in the Shadow of Firm Turnover: ISPs in the 1990s." *International Journal of Industrial Organization.*

Strover, Sharon, 2001, 'Rural Internet Connectivity'. *Telecommunications Policy*, 25(5): 331-47.

Stiroh, Kevin. 2002. Information Technology and The U.S. Productivity Revival: What do the Industry Data Say? *American Economic Review*.

Tedlow, Richard, 2006, *Andy Grove: The Life and Times of an American*, Penguin Group; New York.

Manuel Trajtenberg, 1990, *Economic Analysis of Product Innovation, The Case of CT Scanners*, Harvard University Press; Cambridge, MA.

Triplett, J. E. 1989, "Price and Technological Change in a Capital Good: A Survey of Research on Computers," in *Technology and Capital Formation*, D. W. Jorgenson and R. Landau (eds) MIT Press, Cambridge, MA., pp 127-213.

Varian, Hal 2000. “Versioning information goods.” In Brian Kahin and Hal R. Varian, editors, *Internet Publishing and Beyond.* MIT Press.

Varian, Hal, 2001, “High Technology Industries and Market Structure,” September, Working paper, http://www.ischool.berkeley.edu/~hal/papers/structure.pdf.

Von Burg, Urs, 2001, *The Triumph of Ethernet: Technological Communities and the Battle for the LAN Standard*, Stanford University Press.

Von Hippel, Eric, 2005, “Open Source Software Projects as User Innovation Networks,” in (ed) Joseph Feller, Brian Fitzgerald, Scott A. Hissam, Karim Lakhani, *Perspectives on Free and Open Software,* MIT Press; Cambridge, MA. Pp. 267-278.

Waldrop, Mitchell, 2001, *The Dream Machine: J.C.R. Licklider and the Revolution that Made Computing Personal*, Penguin Books; New York.

West, Joel, and Scott Gallagher, 2006, “Patterns of Innovation in Open Source Software,” in (ed) Hank Chesbrough, Wim Haverbeke, Joel West, *Open Innovation: Researching a New Paradigm,* Oxford University Press; Oxford, pp 82-108.

Figure 1.

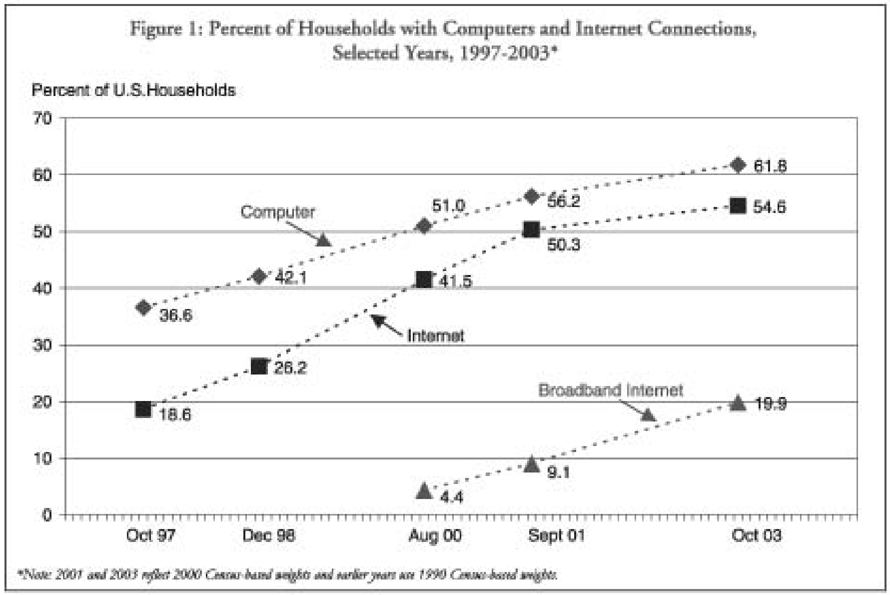

Source: NTIA, 2004.

Figure 2.

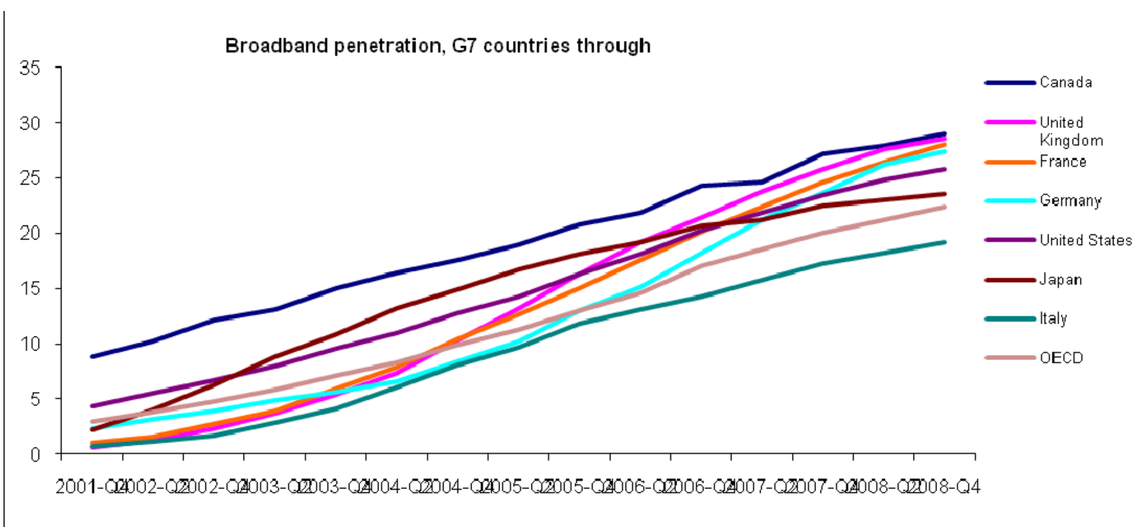

Source: OECD Broadband Portal, http://www.oecd.org/sti/ict/broadband, Table 1i.