



Ubiquitous Clicks and How It All Started

I think the press has a tendency to pick a person and paint them 10 feet tall. In fact, each of us does a little piece and I've done one thing, people add on that and then another. So you get credit for doing the whole damn thing, and that's not so.

—Paul Baran, *after receiving the National Medal of Technology and Innovation*¹

One day I watched my children use the Internet and soon found myself talking with them about the Internet in the same effusive way my immigrant grandparents talked about the wonders of electricity and the magic of transcontinental air flights. My kids just shrugged their shoulders at their father's dramatics and went back to surfing the web and playing online games.

My children cannot imagine a world without the Internet. Clicks are familiar. Hasn't it always been so?

Modern economies frequently change frontier technologies into widely used ones—from the mysterious to the unremarkable. The Internet was once exotic to all but a small set of cognoscenti, but long ago the technology spread to a majority of households and businesses. In the process of becoming ubiquitous it transformed how we work and live—changing how consumers behave, and altering how firms provide products and services.

¹The quote comes from Cassidy (2011).

My children are not alone in their shrugs. Most adults today could not say where the Internet came from and how its spread caused so many transformative changes. Not only does that hold for many educated adults, it also holds for many of society's thought leaders. I have met many educated economists who know a great deal about many technologies, and yet they lack any view about the general economic lessons the Internet's spread illustrates. I have met many perceptive legal scholars and policy analysts with a similar gap in their understanding.

I have also met many who do not shrug, who are curious, especially among those too young to have lived through the relevant events. Can a knowledgeable observer explain how and why the Internet deployed as it did? To what do we attribute its impact? Can those recent events yield insights that help understand technology in the future? They do not know where to go to answer their curiosity. Many key events were genuinely complex, and, at times, involved a vast ensemble of participants with distinct motives and alternative points of view. It is not obvious where to start.

This book addresses this curiosity and points it to the deeper mystery behind the surface of events. The Internet's deployment is commonly held responsible for an economic boom. Along with that boom, and in less than one generation, the names of the leading suppliers of communications changed. So too did the predominant view for forecasting how the underlying technology in communications would evolve. To the common eye all these changes occurred in less than a decade, which is extraordinarily fast by historical standards. Any of these would be rare to see in any industry. The combination—economic boom, change in leadership, alteration of the common forecast, and rapid change—is rarer still in the history of modern capitalism. These are typically associated with only the most transformative technologies, such as the steam engine, the railroad, electricity, indoor plumbing, and the automobile. Such a combination of events merits an explanation in its own right, because the history of capitalism suggests this should not happen often, if at all.

Existing explanations leave a gap, however. While many rich and wonderful histories have been written about the invention of the Internet, most focus on just invention. That yields an answer rich in technical details and incomplete in perspective. It diminishes the role of markets and government policies for markets. Writing about policy addresses some of that gap but tends to stress legal issues, regulatory debates, and changes

in court decisions. It overlooks economic incentives and the behavior policy induces from suppliers and users.

To say it broadly, comparatively less writing focuses on explaining the Internet's innovation and commercialization. Innovation is the act of turning invention into something useful, while commercialization translates innovations into valuable products and services. Innovation and commercialization must connect to each other because both involve market activities, such as building the production and distribution processes to deliver a new service to customers. That summarizes the motivation behind this book and its outlook. It is not possible to explain the deeply surprising and unique aspects of these events—economic boom, change in leadership, alteration of the common forecast, and rapid change—only with technology and policy analysis. It also requires understanding how innovation commercialized, that is, how innovations became valuable as the Internet evolved within commercial markets.

The commercialization of the Internet merits attention for a related reason—it illuminates important relationships between transformation in industrial structure and innovation. How can innovation have such a transformative role in the restructuring of industries? Specifically, replacing one economic structure with another is often called the process of “creative destruction,” and a crucial question of this book might be rephrased as “What role does innovation and commercialization play in creative destruction?”

This rephrased question needs a little refining because the phrase “creative destruction” has become an overused colloquial expression. To get clear on its meaning, return to its most famous user, Joseph Schumpeter. In his 1942 book *Capitalism, Socialism, and Democracy* he describes creative destruction as an inherent property of market-based capitalist systems.

Capitalism, then, is by nature a form or method of economic change and not only never is but never can be stationary. . . . The fundamental impulse that sets and keeps the capitalist engine in motion comes from the new consumers' goods, the new methods of production or transportation, the new markets, the new forms of industrial organization that capitalist enterprise creates. . . . This process of Creative Destruction is the essential fact about capitalism.

Schumpeter argues that little is permanent in market economies, and such impermanence serves a useful purpose in the face of large dominant

firms. The market value of dominance will motivate new firms to aspire to reach a similar position, principally through creating and promulgating innovations. This activity motivates incumbent firms to fend off such entrants, also by innovating and offering those innovations in commercial markets. Altogether, says Schumpeter, more innovation results from the tournament to become a dominant monopolist, and, relatedly, from defending a lucrative monopoly position by innovating faster than those who threaten it.

Schumpeter's conceptualization directs attention at the role of contests to become monopolies in innovative activities, but it leaves open critical questions about how to best organize such a tournament in a market-based system. Many writers have considered that topic in a variety of technologies and time periods, and this book does as well, and does so in the case of the commercial Internet in the 1990s. Now a core motivation for the book can be rephrased into a seemingly simple retrospective question: What aspects of market structure played an essential role in fostering the growth and development of creative destruction during the deployment of the Internet?

The book's answer to this question focuses primarily on the experience in the United States. Why do this when the commercial Internet has reached a global scale in operation and in final service markets? In part, the US experience generates a coherent analysis and a surprising narrative. Many of the most important early innovations took place in the United States under the stewardship of the US government. The transition out of US government stewardship had a profound impact on how the Internet shaped commercial markets. The Internet also grew rapidly due to some unique structural features of the market for communications in the United States, as well as some unique environmental factors supporting entrepreneurial ventures. Many of these economic factors are not widely appreciated, and focusing on them yields novel insights.

Concentrating on the US experience also directs critical analysis and skepticism toward one the most controversial and important lessons from this experience—the role of government policy in fostering economic growth through encouraging innovation. No expert doubts that the US government played an important role in events—as buyer, lawmaker, regulator, and booster—but there is considerable debate about whether its role was salutary and purposeful. Did government policy and action nur-

ture a positive outcome, or did the key features of the outcome arise in spite of the government role? If government actions helped, what mechanism brought about positive contributions, and why, and when can a government use such policies again? If government policy was a hindrance, what should a government avoid doing? This book aims to identify and sharpen the answers to these questions with a comprehensive examination of the experience. The focus on innovation and commercialization will lead to policy lessons that can carry over to commercialization of other technologies.

The focus on the United States also was a pragmatic choice. There were many networks around the world, and some of them, such as Minitel in France, achieved large-scale use and met some benchmarks of success. Some attempts at standardization of networking, such as the efforts put into open systems interconnection (OSI), did not realize their aspirations before folding some of their successes into the overwhelming advance of the Internet. It is certainly plausible that comparison between countries—for example, between the economic experiences within the United States, Japan, and many European countries—could identify additional factors that shaped outcomes. The required detail is simply beyond my grasp, however. Comparative questions will remain for another analyst to explore and explain.

One of the other goals of this book is to identify the grain of truth in well-known economic myths, and to dismiss the falsehoods. Economic myths are misleading economic metaphors. Usually an economic myth is based on a misreading of a short story or aphorism, and it points away from the key lessons of an event. Many economic myths purport to explain how the Internet developed in the United States, and some of them will be familiar to most readers merely from reading the news. Many economic myths about the Internet have not been confronted by scrutiny. This book must confront many of those myths and correct them.

Perhaps the most pernicious of those myths is *Internet exceptionalism*. This is the belief that the Internet followed its own unique economic rules, having little in common with other important historical episodes. This belief was common in the United States during the late 1990s, voiced with enthusiasm by participants in the dot-com boom, primarily in the entrepreneurial sector. Internet exceptionalism also can be an ideology that overly stresses the role of the unique features of Internet technology in

commercial events. It misattributes most novel economic gains to technical causes, and was commonly espoused by those entrepreneurs familiar with the Internet's technical features. All flavors of Internet exceptionalism relegate economic analysis of commercial behavior and incentives to secondary status and deemphasize or overlook the influence of commercial markets in fostering or discouraging innovation. This book argues that Internet exceptionalism is just plain wrong, that it can be replaced with a coherent and sound economic explanation, and, crucially, that this explanation must serve as the foundation for understanding the broader lessons about the Internet's evolution, for explaining the innovation that took place, and for analyzing the causes behind creative destruction.

Replacing myth with sound economics faces numerous challenges when it comes to the Internet because no single participant experienced the entire event. No single story can narrate its causes. In the 1990s there was no such thing as a typical Internet company or a typical Internet strategy, or a typical user of the Internet or typical application for the Internet. There also was no such thing as an advanced plan for the Internet, and no single organization orchestrated the design, building, and operation of the Internet. The relevant experiences span multiple generations of participants from a varied set of backgrounds—for example, government laboratory managers, university computer science graduate students, Internet service providers and their commercial cousins in bulletin board firms, start-up veterans and founders, and platform software makers, among many other participants.

That requires a book that mixes stories, economic insight, and general lessons. Accordingly, this book employs a framework for each chapter that generates this mix. Each chapter will begin with one story. Each story motivates a broader examination of a specific question or set of questions about commercialization, which the chapter pursues. In the end, each chapter provides a few general economic insights and lessons. The conclusion reviews those insights and shows how the lessons fit together under one framework.

What role do the stories play? This is easier to illustrate than to explain in general terms. Here are three examples of such stories:

- Chapter 2 begins with a story describing the experience of assistant attorney general William F. Baxter. He visited the White House just prior to the divestiture of AT&T. The White House did not inter-

vene in the divestiture, and that nonevent is the first of many to motivate inquiries about planning the Internet. Did any single decision maker or executive orchestrate the Internet in a top-down fashion? Generally no, and, as later chapters show, despite some bumps in the road, it turned out well for society.

- Chapter 6 begins with a historical story, the beginning of the California gold rush in 1848. This story motivates analysis of the conditions that produced the temporal concentration of economic activity, which common language labels as a “rush.” Understanding the detail of the story begins the broad inquiry into whether events in 1995 actually resembled a gold rush. It did, but not for too long, and, as later chapters explain, something other than a gold rush explains why the boom sustained itself and eventually crashed.
- Chapter 13 starts with a discussion about the policies that governed National Science Foundation (NSF) funding of research. It shows how the NSF’s flexible policies helped two Stanford professors, Hector Garcia-Molina and Terry Winograd, and their students, among them Larry Page and Sergey Brin. The latter two founded Google with a project that started in their professor’s labs. This story begins a larger narrative about how markets renew themselves with new ideas that run contrary to assumed wisdom. Did the NSF intend to renew the market? No, that was not the direct intent of its funding. However, the flexibility of their funding process helped indirectly, because it raised the chances that the research would be relevant. In the end, society benefited from that flexibility.

These are just three examples among many, and they illustrate the general challenge in a book with these goals. Commercialization of any major technology, including the Internet, does not involve one individual in one location in one time period. Each participant’s story draws on details that reside in different organizations and distinct time periods. Each story illuminates distinct economic forces, decisions, and policies. In each case the story drives toward understanding the sound economic reasoning illustrated by events. The larger narrative is comprised of many of these stories and insights because no single insight or story could or should address the key questions motivating the book. This topic requires an economic narrative with wide scope that integrates many insights.

Innovation from the Edges

This book is more than stories; it describes general economic principles and analyzes connections between cause and effect using economic analysis. It shows how different factors worked in the same direction or against one another to produce the observed outcome. The book distinguishes among three categories of causes—economic archetypes, government policies, and influential institutions:

- *Economic archetypes* are patterns of economic behavior that reflect economic forces and principles that manifest repeatedly in different episodes. Economic archetypes are not unique to the Internet and have appeared in other markets or time periods.
- *Government policies* shape behavior and outcomes with the force of law or regulatory authority. These policies are either inherited, reflecting the legacy of prior legal or regulatory decisions, or they are chosen for anticipated impact, reflecting the preferences, wisdom, and foolishness of those with legal or regulatory authority.
- *Influential institutions* describe business practices and social norms that shape actions, though not necessarily with economic motive as the primary driver. These norms often exist for sensible reasons, and sometimes they exist for independent reasons that are unconnected to their economic impact.

What is one goal of the book's analysis? Said succinctly, this book identifies the economic archetypes, government policies, and influential institutions that shaped the growth and evolution of the commercial Internet. More specifically, the book asks whether specific economic archetypes, government policy, and influential institutions encouraged or discouraged innovation from the edges. It also analyzes which economic archetypes, government policies, and influential institutions played an essential role in generating a large economic impact from innovation from the edges.

Longtime aficionados of communications policy may recognize the phrase, "innovation from the edges." In its original use it typically describes experiments in communications markets, often using new wireless technologies outside of major urban markets, where the impact is (typically) small but the discretion to experiment is large. This book appropriates the phrase and employs a more expansive meaning. For the purposes

of this book, *innovation from the edges* describes innovation being commercialized by suppliers who lacked power in the old market structure, who the central firms regarded as peripheral participants in the supply of services, and who perceived economic opportunities outside of the prevailing view. This definition embeds three related interpretations—stressing place, power, or perceptions.

Consider the interpretation around “place.” An *edge* can refer to a distant *place*, an area far from a central location, such as the edge of town. *Innovation from the edges* refers to innovations coming from a dispersed set of places, and each place is located away from a central location. In that sense *innovation from the edges* refers to a pervasive feature of the Internet historical experience: many key innovations for the Internet came from contributors in a widely dispersed set of places. Those innovations were added to, or competed against, the contributions that came from the central locations. The book asks: How and why did some innovations come from places outside the center, and why was that largely a positive experience?

There is another way to interpret *innovation from the edges*. It brings to mind “power” structures in markets. For example, in many markets a large powerful firm inhabits the leadership position, while many less powerful firms live in its shadow in the periphery. In this sense *edge* refers to less-powerful peripheral supplier and their customers, and accordingly, *innovation from the edges* describes innovation arising from firms and users that leading firms perceived to be on the periphery. That sense also fits a pervasive feature of the historical experience for the Internet: many of the key innovations came from the seemingly less powerful and more peripheral participants—research laboratories, bulletin board firms, start-ups of all stripes, consultancies, and iconoclasts working with distinct outlooks. The book asks: How and why did so many seemingly powerless players have such a large influence on so many commercial episodes, and why did that have such a positive influence on outcomes?

A third interpretation of *innovation from the edges* refers to “perceptions.” The sources of information reflect the edges of perception at the leading firms. In this sense, the *edge* refers to a place beyond the horizon of a vision, a perception that fell outside of known forecasts or prevailing points of view. That captures one other pervasive feature of the historical experience: many of the key innovations fell outside known forecasts and predictions and were unanticipated by established firms in computing and

communications. The book asks: How and why did the Internet's deployment lead many firms to reconsider their perceptions about how the Internet could contribute to the generation of market value, and how did changing perceptions shape outcomes?

Why did innovation from the edges play such a large role? Many chapters will stress the importance of uncertainty over the level of value the Internet could create, and how it could be created. To say it concisely, there often was genuine disagreement among experts about the sources of value in Internet markets. The situation supported different opinions about the strategic actions for firms to take. Not every participant perceived the value of the opportunities, or possessed the same assets for taking advantage of the opportunities. Decentralized decision was far better at enabling exploratory actions than a central decision-making process, often because it enabled actions by entrepreneurial actors with distinct points of view.

As it turned out in the United States in this time, innovation from the edges also played a large role because policy allowed it to. Many innovative participants who came from the edges were unrestrained by barriers or frictions, and policy played a role in fostering the low barriers and frictions (in a sense explained in the book). These aspects of the situation are complex or involve obscure detail, so they are often underappreciated. Accordingly, the book spends some extra effort providing details about policy and explaining their pertinence to economic behavior and outcomes.

Schumpeter would have forecast the narrative arc of this book: mainstream firms who had traditionally assumed leadership roles were threatened by innovation from others, and they reacted by innovating. The book also argues that much of this innovation would not have occurred in the absence of innovation from the edges. By fostering creative destruction, innovation from the edges created value much faster and with greater success than any single organization ever could have.

The Layout of the Narrative

The book makes one additional important compromise: it starts its narrative in the late 1980s when the Internet began to "privatize." Privatization involved the change in the ownership status of public assets associated with operating the Internet, changing them into private assets. Privatization involved a complex process, withdrawing the government from own-

ership and management of the Internet backbone, and substituting private or collective not-for-profit organizations in every instance. The earliest two chapters describe this process.

There is a reason to start then, and it goes to the core of this book's novelty. First, while it is always possible to go further back in time, for this book's purposes most of the important commercial events take place after privatization, primarily in the 1990s. Second, the book focuses on something only peripherally explored by others—namely, the links between innovation, privatization, and the birth of the commercial Internet. Again, the key events for that purpose occur after privatization. Said another way, this book has little new insight to offer about events in the 1970s and 1980s, prior to privatization. Many scholars and commentators have covered these events, as the footnotes and references will document. The book develops plenty of novel insight about later events, and that is where the book focuses its attention.

That does not mean the book can ignore events prior to privatization. Privatization occurs after the Internet already has incubated under government stewardship for a considerable time in the 1970s and 1980s. At first a military research organization had sole responsibility for managing the precursors to the Internet. A new era began in 1985, when the National Science Foundation accepted responsibility for managing the aspect of the Internet that supported research throughout universities. Privatization began during this second era, around 1989, and eventually brought about the end of government stewardship.

For readers not steeped in the history of the Internet, it will be challenging to appreciate the narrative without understanding a few events prior to privatization. Chapter 2 must recall some of these events. In addition, the end of this introduction presents a brief explanation of events prior to privatization—just enough to meet the needs of a novice to this topic. (A reader who knows the history of the Internet's invention can skip this section and go to chapter 2.)

The logic about where to start the book also determines where to end it—namely, around 2003. The book focuses on a number of events shaped directly or indirectly by privatization. That encompasses many of the events affiliated with the economic boom of the late 1990s. It does not include many events after the dot-com bust, but it does include some important events whose origins can be traced indirectly to privatization, such as, the birth of wireless Internet access, commonly called Wi-Fi, which has

deep roots dating to policy decisions in the 1980s. It also covers the rise of search supported by advertising, as founded by Google, which also traces its origins to research conducting just as privatization got underway.

Now I will describe the organization of the book. The book collects its chapters into three distinct sections that cover partially overlapping sequences of related events. The first section contains four chapters. It focuses on understanding how privatization occurred and why it had such a large impact on commercial markets. The second section contains five chapters. It focuses on understanding why so many disparate participants seem to act in concert after privatization, all investing at the same time, fostering a boom in economic activity. The third section contains four chapters. These chapters highlight how and why many participants in the economy did or did not undertake “exploratory” actions, creating and developing new commercial services embedding Internet technology.

Begin with the first group of chapters. Seen from the perspective of participants in 1989, the Internet had accumulated many of the pieces of the modern Internet. For example, the infrastructure had been successfully engineered to support a packet switching network, and it daily handled a large volume of electronic mail and file transfers, and demand appeared poised to continue to grow. The most vexing issues about privatization were not technical, but institutional, legal, and commercial in nature. The first group of chapters is called “The Transition,” and it is comprised of chapters 2 through 5. These chapters identify factors in privatization that created innovation from the edges. Table 1.1 presents a chronology of the notable events touched on by this group of chapters.

As chapter 2 describes, the Internet should not be understood solely as a network technology. Even prior to privatization, groups met regularly to facilitate the Internet’s growth and development. While no single administrator controlled or planned the Internet, the outlines for what became the Internet’s “governance” already existed in this era. That insight helps understand what had to change to accommodate privatization, which is the topic of the next chapter.

Chapter 3 focuses on the privatization process itself and highlights several key policy decisions that shaped the evolution of the Internet. It highlights the decisions that turned out to be crucial for later commercial experience, such as NSF’s decision to set up the commercial Internet as a competitive market and not a monopoly. Although that might seem like an obvious policy

TABLE 1.1. Chronology: The Transition
Selected notable events from chapters 2, 3, 4, and 5

Year	Chapter	Notable Event
1990	2, 3	NSF conducts conversations about privatization
	3	PSINet and UUNET begin first full year as private firms
1991	3	High Performance Computing Act of 1991 passed
	4	Tim Berners-Lee downloads code for web to shareware sites
	3	Commercial Internet eXchange (CIX) founded
1992	3	Network Solutions takes control of domain name system
	3	Rick Boucher sponsors a bill to amend NSF charter
	3	Internet Society founded and IETF becomes part of it
1993	3	Final NSF plan for privatization emerges, and NSF solicits bids
	4	Mosaic browser made for Unix and Windows OS
	5	Earliest ads for ISPs appear in <i>Boardwatch Magazine</i>
1994	4	Founding of the World Wide Web Consortium
	4	Mosaic Communications Company (MCC) founded
	4	MCC changes name to Netscape, and releases a beta browser
1995	4	Apache formed from different versions of NCSA HTTPd server
	3	NSFNET shutdown, and Internet backbone privatized
	5	Netscape IPO and Windows 95 launched in same month
1996	5	Congress passes the 1996 Telecommunications Act
	5	More than 2,000 ISPs advertise in <i>Boardwatch Magazine</i>

choice in retrospect, the NSF did not reach it by a straight path. The chapter explains how that happened, and begins to explain why the consequences from privatization turned out to be so difficult for contemporaries to forecast.

After commercialization of the Internet many participants online tended to treat the Internet and the web as synonymous, even though their blending obscures the origins and economic functions of the network and the software layered on top of it. The Internet refers to the vast networking

infrastructure that connects computers, while the World Wide Web became a layer of software that enables browser-based applications over the Internet. Chapter 4 considers how a number of policies shaped the transfer of the web's technology out of universities and into commercial markets. Once again, these events were crucial for later commercial outcomes and are often unappreciated.

The Internet did not grow in an isolated research lab, and it did not privatize in a commercial setting that started from scratch. Chapter 5 puts the spotlight on the eclectic and entrepreneurial communities outside of the government who would play a crucial role in building commercial Internet service providers, or ISPs. ISPs provided access to the Internet to homes and businesses in exchange for a fee. The ISP market was not particularly large until after privatization, and this chapter helps understand where many of the entrepreneurs came from. Many came from the bulletin-board industry. This chapter also focuses on understanding why—a short time after privatization—their presence fostered innovation from the edges.

The book then turns to explaining how the Internet “gold rush” emerged—when investment, adoption, and new formation of firms all grew at the same time. It concludes by showing that the metaphor of the “gold rush” both informs and misleads. The metaphor describes only an early part of the experience, while later chapters show that other metaphors provide a more satisfying economic explanation for why growth persisted. A deeper explanation analyzes the new commercial value chain for services related to the commercial Internet. A value chain is a set of interrelated activities that together produce a final product of value greater than the incremental value of each part. All mature industries have a value chain. The next section, chapters 6 through 10, examines events in the middle of the 1990s, when this value chain grew. These chapters compose the section “The Blossoming.” Table 1.2 provides a chronology of notable events from this group of chapters.

The first chapter in this group grapples with the economics of a gold rush. Chapter 6 explains the relationships among privatization, the creation of the web, and the timing and bunching of investment at one time. The creation of the first commercial browser, in particular, resulted in a catalyst for action by many private actors. Why? It was a catalyst, because it was a working prototype of a commercial product that showed how to deliver valuable functionality to users. While plenty of technical histories

TABLE 1.2. Chronology: The Blossoming
Selected notable events from chapters 6, 7, 8, 9, 10

Year	Chapter	Notable Event
1992	7	David Clark speaks of “Rough Consensus and Running Code”
	7	Internet Society founded and IETF becomes part of it
	7	Tim Berners-Lee first visits the IETF to standardize the web
1993	8	Louis Gerstner hired as CEO at IBM
	7	CERN renounces ownership rights to World Wide Web code
	8	Earliest ads for ISPs appear in <i>Boardwatch Magazine</i>
1994	6	Vermeeer founded, begins work on web-authoring tools
	7	Tim Berners-Lee founds the World Wide Web Consortium
	6	Brad Silverberg organizes team at Microsoft to examine web
1995	6	Gates circulates the memo, “The Internet Tidal Wave”
	6, 7	Netscape IPO and the launch of Windows 95
	9	HoTMaiL founded, and “viral marketing” is invented
1996	8	Microsoft offers Internet Explorer at a price of zero
	8	AT&T WorldNet sold at \$19.95 for unlimited service
	8	AOL implements all-you-can-eat pricing
1997	8	56K modems first introduced
	10	Tiered structure emerges among Internet data carriers
	9	Netscape and Microsoft reach parity in browser features
1998	10	WorldCom merges with MCI, spins off backbone assets
	8	Over 65,000 phone numbers available for dial-up ISPs
1999	9	Dot-com boom reaches greatest height
	10	WorldCom proposed merger with Sprint is called off
2000	8	<i>Boardwatch Magazine</i> records over 7,000 ISPs
	9	Internet adoption nears saturation at medium/large businesses

have made such an observation, this chapter places an emphasis on the underappreciated role of commercial markets. It stresses what firms could have done had they known how to deliver value. How did value get delivered to users? Technology alone could not do it. Firms also had to observe a working prototype of a functioning and viable business.

Chapter 7 contrasts the commercial operations of the Internet with the commercial operation of the personal computer market. In 1995, the value chain for the commercial Internet was quite young and still undergoing dramatic change, while the PC market was almost two decades old and dominated by Microsoft and Intel. Despite overlapping in some respects, the two value chains differed in many subtle ways, and those differences will help explain how and why the process of creative destruction took root.

Even after a new value chain comes into existence, its existence alone does not explain why it delivers value. Chapters 8, 9, and 10 explain the economic growth affiliated with building the commercial Internet and web. These chapters analyze how value was created in different sectors—Internet service providers, households, business users, consultants, and other suppliers in established and entrepreneurial firms. These three chapters provide many illustrations of private actors adapting the Internet to their needs in specific market and organizational circumstances, arguing that such adaptation was crucial for economic growth.

These chapters stress how investment in one sector complemented the other. That illustrates an important observation about economic growth in this period: the creation of value in one sector reinforced its creation in another, raising incentives to invest in the Internet. In this instance, such reinforcement generated a “network effect,” in which one decision maker’s participation in the Internet economy raised the value to another’s participation, and it operated at an economy-wide level. It also played out over time as “a virtuous cycle,” in which investment by one actor encouraged investment by another at a later time, and on and on in a chain across multiple sectors. That pattern also provided strong incentives for impatient investment behavior. That is an important economic explanation for why firms continued to invest at the same time in the late 1990s, and it is distinct from the factors that catalyzed investment in the middle of the 1990s.

These chapters explain why the Internet boom may have started as a gold rush, and how it was sustained by a virtuous cycle. That explanation

also will play a role in the dot-com bust, when the virtuous cycle came to a halt.

Exploratory behavior arises in all the chapters that compose the last section, titled “Exploration and Renewal.” Chapters 11 through 14 largely cover events in the late 1990s and the beginning of the new millennium, until the dot-com bust led to a decline in investment and entrepreneurial activity. In other words, these four chapters stress several different exploratory episodes unleashed by innovation from the edges during the first wave of investment. Table 1.3 provides a chronology of this group of chapters.

Chapter 11 focuses on how Microsoft reacted to the new prevailing view about the prospects for the Internet. It perceived a commercial threat and acted to redirect behavior unleashed by innovation from the edges. This event became popularly known as the “browser wars,” and it shaped the commercial Internet for many years. The chapter focuses on understanding the strategic behavior of Bill Gates, and why he sought to discourage Netscape and its business partners. The analysis leads to several insights about how and why leading firms do and do not have incentives to encourage innovation from the edges.

Internet exceptionalism comes in for scrutiny in chapter 12. This chapter analyzes how this ideology shaped behavior of many participants within financial markets and distorted the growth of electronic commerce. Said simply, this chapter identifies many of the factors that distorted the economics of what was popularly called the dot-com boom and bust. The chapter places particular emphasis on impatient exploratory investment, and how financial markets encouraged those distortions. The chapter also explains why these distortions came to an end. Some of the end was inevitable—a by-product of overshooting in the virtuous cycle. Some of it resulted from a distortion, a product of an Internet exceptionalism allowed to run rampant.

Chapter 13 describes one mechanism for renewal of growth after the dot-com bust. It examines the creation of Google, which arose from innovative software created in a university with funding from the NSF. It is not widely appreciated that these innovations originated from government-funded research, so the chapter begins by focusing on the policies for transferring the results into private markets. The chapter then examines how Google continued to explore new features in a commercial setting,

TABLE 1.3. Chronology: Exploration and Renewal
Selected notable events from chapters 11, 12, 13, 14

Year	Chapters	Notable Events
1994	13	Lou Montulli invents the cookie at Netscape
	13	Sergey Brin begins his graduate studies
1995	11	Bill Gates writes "Internet Tidal Wave"
	11	Netscape IPO and Windows 95 launched in same month
	13	Larry Page begins his graduate studies
1996	11	Microsoft begins pressuring partners not to support Netscape
	12	Greenspan makes speech about "Irrational Exuberance"
	12	Wave of new entrants marks start of dot-com boom
1997	14	FCC issues final draft of Part-15 rules for spectrum
	11	Jobs makes deal so IE becomes default browser for Apple
	11	Netscape and Microsoft reach near parity in browser features
1998	11	Senate hearings about Microsoft
	13	Google founded by Larry Page and Sergey Brin
	11	Netscape coalition collapses, AOL eventually buys Netscape
1999	12	Dot-com boom reaches greatest height
	12	Telecom meltdown begins after rule change for CLECs
	14	IEEE committee 802.11 issues design a & b, labeled "Wi-Fi"
2000	11	Judge Jackson issues judgment that finds against Microsoft
	12	NASDAQ reaches its peak in stock valuations of dot-coms
2001	14	Wi-Fi becomes available on Windows-based systems
	12	PSINet declares bankruptcy
	12	9/11 terrorist attacks on World Trade Center
	11	DOJ settles with Microsoft
2002	12	Internal accountant discovered fraud at WorldCom
	13	Google scales quality-weighted second price auction
	12	Economic decline reaches its nadir
2003	13	Google launches AdSense
	14	Intel launches Centrino

developing both new technologies and new business processes, using automated auctions to sell advertising. It developed a novel approach for keyword search auctions that supported advertising focused on the needs of users.

The final chapter of the narrative, chapter 14, examines the creation of the market for wireless Internet access, which added an important functional capability to the Internet. This innovation was more difficult to create than commonly appreciated, in part because the rules for spectrum use emerged out of a long and complicated policy debate. Like prior chapters, the analysis focuses on the mix of technology and commercial incentives that led firms to explore developing new technical features and new business processes. In this instance, Wi-Fi emerged from a mix of collective private action from a standards committee and entrepreneurial action from several different firms. The chapter ends with general lessons about how firms explored innovative activities in commercial markets, and why some market settings encouraged or discouraged such innovative exploration.

Such an extensive narrative needs a summary and a synthesis of the main lessons. Chapter 15 provides these, and it is called “Enabling Innovation from the Edges.” First it provides a synthetic summary of the book’s story for how and why the commercial Internet grew and evolved. It next provides the summary to the set of questions that motivate the book: What economic archetypes, government policies, and influential institutions encouraged or discouraged innovation from the edges, and why? After reviewing the list of influential factors, the conclusion highlights why innovation from the edges emerged. It also summarizes why rapid market-based learning played a big role in magnifying the impact of innovation from the edges.

The last chapter also offers a surprising set of conclusions. The list of important causes is long and varied. It is not surprising that no single economic archetype or policy alone accounts for such a broad array of events, but this raises a curious observation. Despite the absence of any large coordinating government planner, many independent factors tilted in the same direction, reinforcing one another at a market-wide level. How could that have happened? It was as if innovation from the edges arose from either a vast web of coordinated action or an impressive conspiracy of propitious accidents, and neither is plausible. The book ends with a set of observations about how institutions tended to push different events toward similar

policies and, hence, toward the operation of economic archetypes that generated similar types of outcomes. This helps explain why events appeared to be coordinated, even when most were not.

The book also offers lessons for how society can try to rely on systematic policies—and more than mere luck—to produce similar results in the future. It suggests how policy could be tailored to shape events in the future. Most readers already can anticipate the big theme that runs throughout this part of the conclusion. Government policy can fruitfully address many crucial open issues about a firm's action or a government's policies by organizing the inquiry around a broad question: does an action encourage or discourage innovation from the edges? Addressing that key question provides the crucial insight for understanding why outcomes acquired their specific economic characteristics and contours.

How It All Started

Many key inventions for the Internet began with military funding. Yet the military's sponsorship is easy to misunderstand. It may be tempting to compare the Internet to historically archetypical big invention sponsored by hierarchical government organizations, such as the Manhattan and the Apollo projects. These archetypes for developing technical breakthroughs are not good models for understanding what happened during the military's sponsorship. The Internet was not a single urgent project in a single lab devoted to engineering a single object.

This early history is not widely appreciated. That is the point of this section—to introduce a novice reader to a few known and crucial details about the origins of the Internet.

Another popular myth about the Internet's military origins also interferes with understanding its invention. According to this popular myth, the government developed the Internet in order to survive nuclear war. There is a grain of truth to this myth because Paul Baran's theoretical work at Rand was motivated by a research quest, to design networks that remained robust to damage from war.² Frontier computing also had played an important role in military applications, and this was widely under-

²Although his insights were widely unappreciated at the time, Baran later became well known for this early vision. For more in-depth analyses of Baran's contribution, see, e.g., Abbate (1999), Waldrop (2001), or Norberg, O'Neill, and Freedman (1996).

stood inside the military.³ However, this myth points in misleading directions. Surviving nuclear war was, at most, one of many motivations for the funding for the invention of what became the Internet. More concretely, it had little influence on the actual inventive activity of the researchers who did the inventing.⁴

What should replace the popular myths?⁵ The military did not take action in an isolated research laboratory. Rather, the military funded several inventions, and so did other parts of the government, and so did private industry. Sometimes these inventions complemented one another, and occasionally they substituted for one another. There were (occasionally) porous boundaries between the communities who invented for the military and for private industry, so lessons learned in one domain (eventually) spilled into another. There also were multiple efforts outside the United States, most prominent among them was OSI, which later chapters will discuss, and the efforts competed with one another, and imitated one another.⁶

The Internet's early development, largely but not wholly located in the United States, fit into an economic archetype often called "collective invention." Collective invention "is a process in which improvements or experimental findings about a production process or tool are regularly shared."⁷ What we today call the Internet began as a series of loosely connected engineering projects, with military funding supporting some of those projects. Those projects eventually involved a vastly dispersed set of technically adept participants with a shared interest in the project, but otherwise heterogeneous needs and outlooks. The Internet developed slowly throughout the 1970s and 1980s and accumulated capabilities over time from an enormous number of contributors. Researchers with government sponsorship contributed some of the primary innovation, while plenty were borrowed from the active private sector.⁸

Five partially overlapping groups had major roles in shaping the attributes of the Internet that commercialized in the 1990s. Each group valued

³See, e.g., Flamm (1988), or Edwards (1997).

⁴Again, see e.g., Abbate (1999), Waldrop (2001), or Norberg, O'Neill, and Freedman (1996).

⁵For an extensive development of this view, see Greenstein (2010b).

⁶See Russell (2014).

⁷See, e.g., Allen (1983), Meyer (2003).

⁸For an extensive development of this view, see, e.g., Campbell-Kelly and Garcia-Swartz (2013).

distinct dimensions of functionality, and each altered the accumulation of innovative features over time.

The first two communities were the primary decision makers at funding agencies—the Department of Defense (DOD) and the National Science Foundation. The remaining three, no less important, were programmers/developers/inventors, administrators, and application users. Many were funded by the government agencies and given considerable discretion. Others became participants over time and added their own contributions within their own budgetary limitations. An extensive group of inventors also remained active outside of government circles, and, as this book will discuss in several chapters, had a large influence on how the Internet commercialized after privatization.

The earliest funding for the Internet took a form unlike a traditional military procurement project. The Defense Department organized the project in a subagency called DARPA (Department of Advanced Research Projects Agency), which focused on fostering pathbreaking invention.⁹ What became the Internet was but one of many projects on the frontiers of computer science funded by a special office within DARPA.¹⁰ The project, building of prototypes for a packet switching data-communications network of networks, pushed the boundaries of network computing at the time.

A packet switching network sends messages between computers. It translates the zeros and ones from one computer into many discrete “packets” of data, each of a fixed size. Large packets are divided into many smaller packets, and those packets are then sent between the computers. The computers sending and receiving the packet use the same processes, or “protocols,” for creating and assembling packets. Each packet reserves room at the beginning for identifiers and other code, put there by the computer sending the message. The computer receiving the packet can use the identifiers and code to reassemble those packets, and put them back together in the right order.

DARPA’s program officers understood from the outset that research in packet switching would represent a technical break with prior precedents. Chapter 2 will go into greater detail about how the operations for this

⁹This organization was originally founded as the Advanced Research Projects Agency, or ARPA, and for the sake of simplicity I will use only one name throughout, DARPA.

¹⁰See Aspray and Williams (1994), Norberg, O’Neill, and Freedman (1996), Edwards (1997), Roland and Shiman (2002), or Russell (2014).

technology required more than merely a few inventions. It is sufficient for this backstory to understand that initially packet switching was a budding theoretic concept for how a network should operate. It had not been implemented in the 1960s and could have been implemented in a variety of ways in the 1970s. DARPA's administrators wanted innovations in the form of ideas, new designs, and new software. The inventive goals were large and ambitious, as well as open ended, and that meant the opportunity could not be addressed by a single organization, or by the insight of one lone genius. The inventors and DARPA administrators also understood the goals broadly and did not presume to know what specific designs and applications would suit their needs.¹¹ They broadly funded pie-in-sky research as well as inventions addressing pragmatic problems with anticipated military applications.

As it turned out, the project accumulated capabilities and became very useful for quite a few purposes, such as transferring files, electronic mail, and other forms of communications. Gradually new purposes were invented for the network, and functionality grew and accumulated. Different audiences perceived that open-ended opportunity in different ways and added distinct inventions to the existing packet switching network.

Military needs also overlapped with civilian needs, and that facilitated moving this technology out of the military in the middle of the 1980s. Why? In part, this was an explicit goal. Administrators at DARPA desired that all of these innovations be portable to military operations in the long run. Such an outlook was required under a statute called the Mansfield Amendment (stipulating that Department of Defense funding be relevant to military's mission).¹² Many of the computers used in the military came from civilian suppliers, so such a pragmatic goal inevitably oriented the project toward the similar problems computer users outside the military experienced.

¹¹David Clark, private communications, September 2008.

¹²Norberg, O'Neill, and Freedman (1996) stress that DARPA's funding of packet switching research in the 1960s and 1970s met concerns about whether the funding was relevant to a military mission, as required by the Mansfield Amendment, which was proposed several times, and eventually passed in 1973. The research anticipated enhancing the "command and control" capabilities of commanders increasingly reliant on their computing resources. Flamm (1988) also stresses the overlap in government procurement requirements for computing and civilian needs. For a large class of activities, making progress on a military problem overlapped with progress on a similar or related civilian problem.

What issues did the US military face with its own computing facilities and operations?¹³ For one, as hinted by the earlier references to nuclear war, the military sought a robust design for a communications network, and the potential value of robustness was self-evident. Keeping communications functioning in spite of a blown or cut line has military value in hostile battlefield conditions. In theory, an inexpensive packet switching network could do this because the path taken by the packets did not matter, and, in principle, could follow many different routes. Thus if one route was damaged and another remained open, then the messages could continue to go through.

An additional technical and pragmatic aspiration also motivated funding. An ideal network could facilitate the movement of data between distant computer systems, and with as little human intervention as possible in the intermediary points of the network. A packet switching network could cover vast geographic distances, which could support the sharing of expensive computing resources between faraway places, and without the use of operators at switches. That too had self-evident military value, as it would for any large computer user. Coordinating the exchange, combination, and filtering of data between computer systems in different locations generated numerous potential gains for operations.

Several prototypes for this packet switching network were engineered with DARPA's funding. With additional funding, these innovative designs turned into a prototype of an operating network, operated by managers from Bolt, Beranek and Newman (BBN), a research contractor subcontracting through DARPA. A number of researchers and their students became familiar with its principles. The network grew from this unusual origin, covering more locations and more research laboratories. Eventually the system became reliable and could exchange data between computing systems without frequent human intervention. Once again, such automation had value inside the military as well as outside of it.

By the early 1980s the network had value to a community of researchers with few connections to the military research projects, and with even fewer to military operations. A considerable community of hobbyists and

¹³There has been considerable writing on this question, and this summary skims the surface. See, e.g., Abbate (1999), Aspray and Williams (1994), Norberg, O'Neill, and Freedman (1996), Edwards (1997), and Roland and Shiman (2002).

commercial firms with interests in frontier information technology also existed outside of the military circles. Yet access to the military network required military backing. The military grew tired of the inconveniences of managing participation among researchers with sporadic military funding, and, for this and other reasons, spun off a part of the network to researchers.

The NSF Era

The second era started when the National Science Foundation (NSF) began managing a network for the benefit of the research community in the United States. It was more than a simple change in management, but it would have taken uncommon prescience to understand how crucial it would be for fostering innovation from the edges at a later date. The consequences are easier to see in retrospect.

The NSF invested in the Internet with several motives, principally among them to stretch the Internet's capabilities as an input into research and higher education. That motivation built on a long history of supporting computer science inside universities.¹⁴ In this case, the NSF sought to use the network for more than just computer scientists, and sought to invest in turning the Internet into a large-scale network for researchers, professors, or students coming from a variety of disciplinary backgrounds. Much investment was aimed at aiding basic tasks, such as sending electronic mail.¹⁵ The NSF also aimed to use the Internet to facilitate connecting with supercomputers, making use of the capacity. Supercomputers were expensive fixed investments with no geographic mobility. Some of the core development, especially for software to make the network operate, continued to be managed by the NSF and by the military.

The NSF's investments began to focus on scaling the network as traffic grew, and early into its stewardship it standardized the protocols in its network. It chose a protocol that the military network also had adopted several years earlier, called TCP/IP, which stood for transmission-control protocol and Internet protocol. TCP defined the protocols for sending and receiving packets. IP defined the design of packets and the information

¹⁴See, e.g., Aspray and Williams (1994).

¹⁵See Abbate (1999). Other factors also shaped investments, as discussed below.

contained in each packet. TCP/IP had been defined many years earlier with DARPA funding, and the military also used it.¹⁶ This was a crucial choice, as the commercial Internet still runs on TCP/IP to this day.

The NSF aimed to build a routine and reliable network infrastructure, making it easy to adapt and spread to every university, community college, and research institute.¹⁷ By 1989 the NSF had rationalized the processes and infrastructure underneath the Internet. That investment gave a wide range of participants—students, faculty, and administrators—a taste for what the TCP/IP packet switching network could do to help them in their work—namely, transmit e-mail, send and receive electronic files, and do both reliably over long distances. By this point virtually all users had begun to shorten “inter-network” to “Internet,” which was a more convenient label.

After a few years under the NSF’s stewardship, the Internet had a different look and feel than it had had in the earlier military era. All participants in the earlier era experimented and explored the frontiers of networking, inventing new capabilities for the nascent network. Some continued to do that in the second era, but many of the new users had no intention of inventing anything, and most had no idea how the entire system operated, nor did they want to know. The Internet was merely something they used in their research and work.

Scale brought with it a new set of managerial questions: could private firms perform the same activity as efficiently, or more cheaply? If they could, how should they be organized? These questions would not—could not—be definitively addressed by the NSF. Several later chapters will explain how these questions were addressed, and why the answer coincided with the emergence of innovation from the edges.

A related pragmatic aspiration—cost reduction—came along with the concern about reliability at large scale. That emerged as a central concern

¹⁶Russell (2014) stresses that DARPA supported scientific inquiry with a robust and questioning conversation, but its sponsor also modulated many forms of dissent and discretion. The US military had a lot of leverage because it could withhold funding, and such coercive tactics proved rather effective in getting TCP/IP adopted. Russell also mentions government representatives could compel TCP/IP compatibility through procurement, as they eventually did by requiring TCP/IP in all forms of Unix sold to the military. Russell stresses that these instruments worked because TCP/IP was not an empty promise. The protocol had been effectively deployed and reflected a clear conception of how it should operate.

¹⁷See, e.g., Frazer (1995), Leiner et al. (2003).

at the end of the 1980s because, in part, the NSF did not have a large or reliable budget for operating the Internet or upgrading it. Every large investment required a large appropriation, and that required congressional approval. Congress was not known for writing blank checks to the NSF.

When he took the job of managing the Internet for the NSF in 1986, Stephen Wolff believed the Internet would be useful for a wide set of potential users in all research communities. An engineer by training, Wolff enjoyed working with the network in his position at a government laboratory, the army's research network at Aberdeen, which had a connection to the network sponsored by DARPA.¹⁸ Wolff's attitude, common to the earlier generation of Internet pioneers, combined idealism—belief that the deployment of the Internet could be transformative—with an engineer's pragmatism. Looking back on his experience, with characteristic laconic description, Wolff thought the NSFNET "looked like a good project with good values."¹⁹

Stephen Wolff did not take the job of managing the Internet at NSF in order to privatize the network. Rather, that proposal emerged from daily experience with the reliable network, and from confronting the budgetary realities of operating the network at NSF. He eventually concluded that private firms probably could handle all the relevant tasks. Wolff made an educated guess that the costs for universities and researchers could be lower if private providers supplied services to both his constituents and private users. The budgets for many networks also could improve if there were multiple source of revenue.²⁰

The proposal for privatization did not emerge overnight. Cautiously at first, Wolff widened the conversation, seeking to figure out how universities could share the infrastructure with private users. In 1989 Wolff began meeting with other stakeholders in the academic Internet and so began a series of conversations about introducing private enterprise into the Internet's operations. Many of these conversations took place at the Harvard

¹⁸That also was near the University of Delaware, the research home to David Farber, one of the key participants in building the network at that time.

¹⁹Stephen Wolff, private communications, July 2008.

²⁰Stephen Wolff, private communications, July 2008.

Kennedy School of Government, organized by Brian Kahin, director of the Harvard Information Infrastructure Project.²¹

The participants left copious records of their perceptions at the time. They generally acknowledged that the research-oriented Internet had matured, moving beyond its “nuts and bolts” stage of development.²² While no serious networking engineer thought the Internet’s technical capabilities had stopped evolving, by this point the Internet had acquired many appealing attributes. It was a large-scale and reliable data communications network with a documented code base upon which many participants could build additional layers of applications. Most of the participants worried about losing those accomplishments if the privatization became neglectful of key aspects of the operations. Most of them did not give any thought to fostering anything related to innovation from the edges.

²¹These meetings are documented in Kahin (1992), which provides a marvelous record of the diversity of thinking at the time. See also November 1990, Request for Comment 1192, titled “Commercialization of the Internet Summary Report,” which provides the initial report of the meeting, accessible at <http://www.rfc-editor.org/rfc/rfc1192.txt>, accessed July 2009.

²²This is the phrase used by Mandelbaum and Manderbaum (1992). They meant that the Internet had moved beyond the stage where just installing it and getting it to work was the primary objective of many IT managers. See also the description of the transition to the T-1 backbone in Frazer (1995), pp. 24–26. Participants faced numerous initial complex technical issues, but overcame them, resulting in a widely recognized technical achievement.