Introduction

MOST INNOVATIONS FAIL. And companies that don't innovate die. This is a book about the process of innovation, about how companies utilize and advance technologies to create new products and services. In today's world, where the only constant is change, the task of managing innovation is vital for companies of every size in every industry. Innovation is vital to sustain and advance companies' current businesses; it is critical to growing new businesses. It is also a very difficult process to manage.

Innovation in the Twenty-First Century: A Tale of Two Models

To paraphrase Charles Dickens, for innovation in this new century, it is the best of times and the worst of times. Industrial technology is advancing our understanding of the natural world at an accelerating rate. In the oldest industry in the world, agriculture, companies are learning to use genetic and genomic technology to make crops more resistant to pests, droughts, and diseases, even as they produce more output per acre. In another ancient industry, retailing, the advances in computing and communications are bringing retailers into closer contact with their customers as well as their suppliers, enabling them to provide more variety with less inventory than ever before. The burgeoning services businesses all benefit from technologies that offer better communications with more capabilities at lower prices. The largest service industry, health care, is experiencing an explosion in our scientific understanding of the forces that create life, with the result being the prospect of longer, healthier lives for us all.

Yet in many ways, it is the worst of times for innovating companies. Many leading companies are having a terrible time sustaining their

internal R&D investments. Take the premier industrial research laboratory of the twentieth century, Bell Labs. Not long ago, Bell Labs would have been a decisive strategic weapon in Lucent's battle with Cisco in the telecommunications equipment market.

Lucent, the telecommunications equipment company created in the breakup of AT&T, enjoyed significant momentum from its spin-off from AT&T in 1996, calling itself "the largest start-up in history." It also inherited the lion's share of Bell Laboratories from the old AT&T, which endowed Lucent with a wealth of research and technology to focus on the telecommunications equipment market. And over the next five years, Lucent enjoyed many victories in the market with its new products. Cisco nevertheless consistently managed to keep up with Lucent, and occasionally got to market ahead of it. Although Bell Labs technologies did create many new products and services for Lucent, Cisco also seemed to introduce many new products and services, despite its lack of anything like the deep research capabilities of Bell Labs.

Though they were direct competitors in a very technologically complex industry, Lucent and Cisco were not innovating in the same manner. Lucent devoted enormous resources to exploring the world of new materials and state-of-the-art components and systems, to come up with fundamental discoveries that could fuel future generations of products and services. Cisco, meanwhile, did practically no internal research of this type.

Instead, Cisco deployed a rather different weapon in the battle for innovation leadership. It scanned the world of start-up companies that were springing up all around it and that were commercializing new products and services. Some of these start-ups were founded by veterans of Lucent, AT&T, or Nortel. These people took the ideas they worked on at these companies and attempted to build companies around them. Sometimes, Cisco would invest in these start-ups. Other times, it simply partnered with them. And more than occasionally, it would later acquire them. In this way, Cisco kept up with the R&D output of perhaps the finest industrial research organization in the world, without doing much internal research of its own.

Lucent's experience with the limits of its research capability is not unique. IBM's research prowess in computing was of no avail against Intel and Microsoft in the personal computer business. Similarly, Nokia has catapulted itself ahead of Motorola, Siemens, and other industrial titans to the forefront of wireless telephony in just twenty years, building on its industrial experience from earlier decades in the low-tech industries of

wood pulp and rubber boots. GE's labs are no longer the powerhouse they once were. Xerox has now formally separated from its famous Palo Alto Research Center. Hewlett-Packard's HP Labs have been broken up between HP and Agilent.

This leads to a number of paradoxes that confront all innovating companies in the early twenty-first century. While ideas abound, internal industrial research is less effective. While innovation is critical, the usual process of managing innovation doesn't seem to work anymore. While ideas and external capital are plentiful, companies struggle to find and finance internal growth opportunities. While industrial R&D spending is high, many worry that we are exhausting the "seed corn" of basic knowledge that will propel technology a generation from now.

Not long ago, internal R&D was viewed as a strategic asset and even a barrier to competitive entry in many industries. Only large companies with significant resources and long-term research programs could compete. Research-based companies like DuPont, Merck, IBM, GE, and AT&T did the most research in their respective industries. And they earned most of the profits as well. Rivals who sought to unseat these firms had to ante up their own resources and create their own labs, if they were to have any chance against these leaders.

These days, the former leading industrial enterprises are finding remarkably strong competition from many newer companies. These newcomers—Intel, Microsoft, Sun, Oracle, Cisco, Genentech, Amgen, Genzyme—conduct little or no basic research on their own. Although they have been very innovative, these companies have innovated with the research discoveries of others. And there is a legion of other, even newer companies waiting to supplant these firms if an opportunity should arise. These latter newcomers are also likely to rely on someone else's discoveries to ascend to leadership.

To make matters worse, some companies that made significant long-term investments in research found that some of the resulting output, however brilliant, wasn't useful to them. They found ways to gracefully exit from the further funding of these projects and moved on to more promising work. Then, to their amazement, some of those abandoned projects later turned into very valuable companies. This was the experience of the Xerox Corporation, for example, with its Palo Alto Research Center (PARC). Numerous valuable computer hardware and software innovations were developed at PARC, but few of them made any money for Xerox and its shareholders.

A Shift in Innovation Paradigms

What accounts for the apparent decline in the innovation capabilities of so many leading companies, at a time when so many promising ideas abound? My research suggests that the way we innovate new ideas and bring them to market is undergoing a fundamental change. In the words of the historian of science Thomas Kuhn, I believe that we are witnessing a "paradigm shift" in how companies commercialize industrial knowledge.[1] I call the old paradigm Closed Innovation. It is a view that says *successful innovation requires control.* Companies must generate their own ideas and then develop them, build them, market them, distribute them, service them, finance them, and support them on their own. This paradigm counsels firms to be strongly self-reliant, because one cannot be sure of the quality, availability, and capability of others' ideas: "If you want something done right, you've got to do it yourself."

The logic that informed Closed Innovation thinking was an internally focused logic. This logic wasn't necessarily written down in any single place, but it was tacitly held to be self-evident as the "right way" to innovate. Here are some of the implicit rules of Closed Innovation:

- We should hire the best and the brightest people, so that the smartest people in our industry work for us.
- In order to bring new products and services to the market, we must discover and develop them ourselves.
- If we discover it ourselves, we will get it to market first.
- The company that gets an innovation to market first will usually win.
- If we lead the industry in making investments in R&D, we will discover the best and the most ideas and will come to lead the market as well.
- We should control our intellectual property, so that our competitors don't profit from our ideas.

The logic of Closed Innovation created a virtuous circle (figure I-1). Companies invested in internal R&D, which led to many breakthrough discoveries. These discoveries enabled those companies to bring new products and services to market, to realize more sales and higher margins

because of these products, and then to reinvest in more internal R&D, which led to further breakthroughs. And because the intellectual property (IP) that arises from this internal R&D is closely guarded, others could not exploit these ideas for their own profit.

For most of the twentieth century, this paradigm worked, and worked well. The German chemicals industry created the central research laboratory, which it used to identify and commercialize a tremendous variety of new products. Thomas Edison created a U.S. version of this laboratory, used it to develop and perfect a number of important breakthroughs, and founded General Electric's famed laboratory. Bell Laboratories discovered amazing physical phenomena and harnessed its discoveries to create the transistor, among its many important achievements. Moreover, the U.S. government created an ad hoc central research laboratory to conduct a crash project on nuclear fission, which led to the development of the atomic bomb.

Figure I-2 depicts this Closed Innovation paradigm for managing R&D. The heavy solid lines show the boundary of the firm. Ideas flow into the firm on the left and flow out to the market on the right. They are screened and filtered during the research process, and the surviving ideas are transferred into development and then taken to market.

In figure I-2, the linkage between research and development is tightly coupled and internally focused. Our extant theories of managing R&D are built on this conception. Examples of this thinking are the stage gate process, the chain link model, and the product development funnel or

FIGURE I-1

The Virtuous Circle

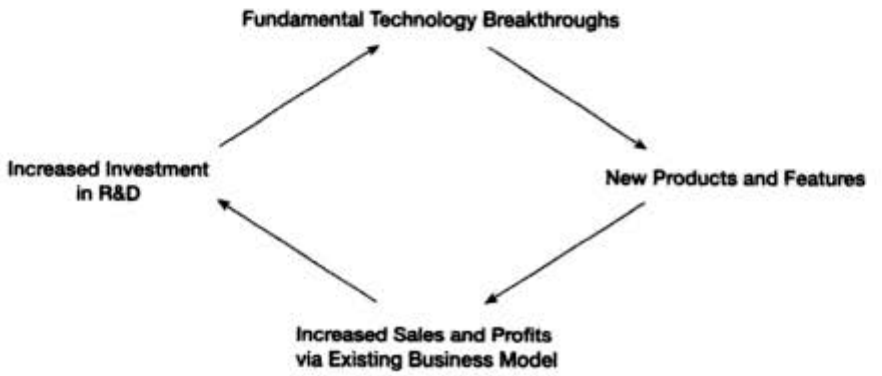

FIGURE I-2

The Closed Paradigm for Managing Industrial R&D

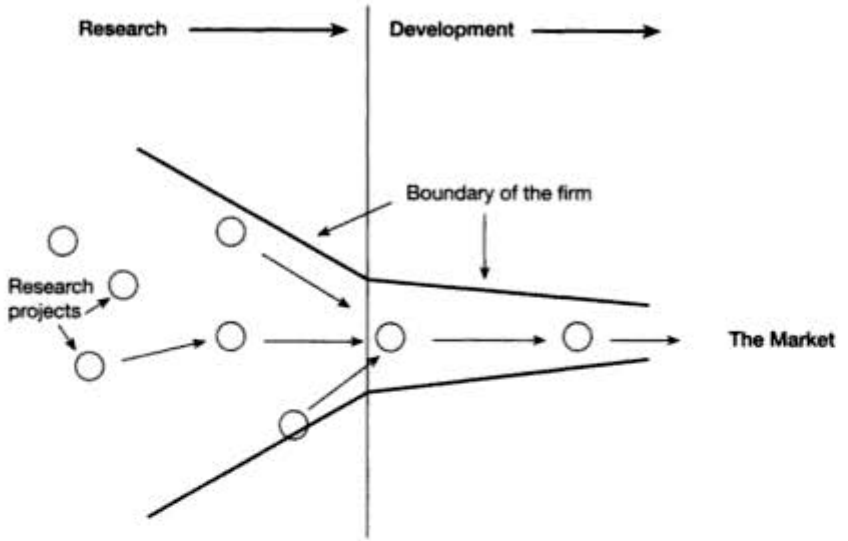

pipeline found in most texts on managing R&D.[2] Projects enter on the left at the beginning, and proceed within the firm until they are shipped to customers on the right of the figure. The process is designed to weed out *false positives*, projects that look initially appealing, but later turn out to be disappointing. The surviving projects, having survived a series of internal screens, hopefully have a greater chance of success in the market.

Erosion Factors That Undermined the Logic of Closed Innovation

In the last years of the twentieth century, though, several factors combined to erode the underpinnings of Closed Innovation. One factor was the growing mobility of highly experienced and skilled people. When people left a company after working there for many years, they took a good deal of that hard-won knowledge with them to their new employer. (The new employer, though, neglected to pay any compensation to the previous employer for that training.) A related erosion factor was the burgeoning amount of college and post-college training that many people obtained. The growing number of such people allowed knowledge to

spill out of the knowledge silos of corporate central research labs to companies of all sizes in many industries. A further factor was the growing presence of private venture capital (VC), which specialized in creating new firms that commercialized external research and converting these firms into growing, valuable companies. Often, these highly capable start-up firms became formidable competitors for the large, established firms that had formerly financed most of the R&D in the industry—the very ideas these new companies fed off of as they competed for industry leadership.

The logic of Closed Innovation was further challenged by the increasingly fast time to market for many products and services, making the shelf life of a particular technology ever shorter. Moreover, increasingly knowledgeable customers and suppliers further challenged the firm's ability to profit from their knowledge silos. And non-U.S. firms became more and more effective competitors as well.

When these erosion factors have impacted an industry, the assumptions and logic that once made Closed Innovation an effective approach no longer applied (figure I-3). When fundamental technology breakthroughs occurred, the scientists and engineers who made these breakthroughs were aware of an outside option that they formerly lacked. If the company that funded these discoveries didn't pursue them in a timely fashion, the scientists and engineers could pursue these breakthroughs on their own—in a new start-up firm. The start-up company would commercialize the breakthroughs. Most often, the company

FIGURE I-3

The Virtuous Circle Broken

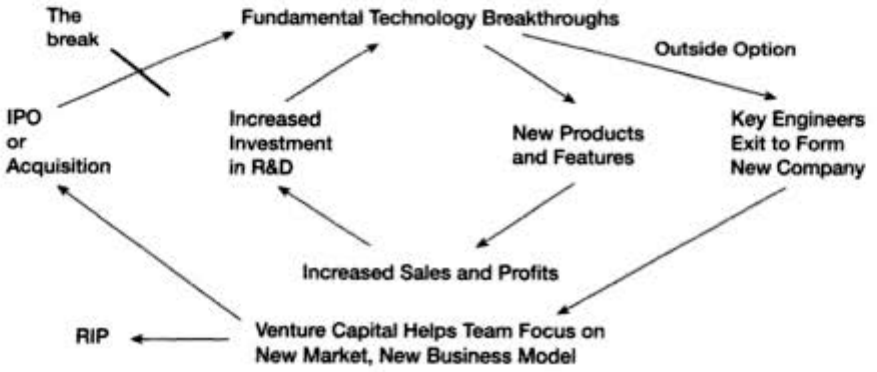

failed (shown in figure I-3 as Rest in Peace [RIP]). But if it became successful, it might achieve an initial public offering (IPO) or be acquired at an attractive price. The successful start-up would generally *not* reinvest in new fundamental discoveries. Like Cisco, it would instead look outside for another external technology to commercialize.

The presence of this outside path broke the virtuous circle. The company that originally funded the breakthrough did not profit from its investment in the R&D that led to the breakthrough. And the company that did profit from the breakthrough generally did not reinvest its proceeds to finance the next generation of discovery-oriented research. This severed link between research and development meant that there would not be another round of investment in basic research to fuel another round of advances.

In situations in which these erosion factors have taken root, Closed Innovation is no longer sustainable. For these situations, a new approach, which I call Open Innovation, is emerging in place of Closed Innovation. Open Innovation is a paradigm that assumes that firms can and should use external ideas as well as internal ideas, and internal and external paths to market, as the firms look to advance their technology. Open Innovation combines internal and external ideas into architectures and systems whose requirements are defined by a business model. The business model utilizes both external and internal ideas to create value, while defining internal mechanisms to claim some portion of that value. Open Innovation assumes that internal ideas can also be taken to market through external channels, outside the current businesses of the firm, to generate additional value. Figure I-4 illustrates this Open Innovation process.

In figure I-4, ideas can still originate from inside the firm's research process, but some of those ideas may seep out of the firm, either in the research stage or later in the development stage. A leading vehicle for this leakage is a start-up company, often staffed with some of the company's own personnel. Other leakage mechanisms include external licensing and departing employees. Ideas can also start outside the firm's own labs and can move inside. As figure I-4 shows, there are a great many potential ideas outside the firm. In figure I-2, the solid lines of the funnel represented the boundary of the firm. In figure I-4, the same lines are now dotted, reflecting the more porous boundary of the firm, the interface between what is done inside the firm and what is accessed from outside the firm.

FIGURE I-4

The Open Innovation Paradigm for Managing Industrial R&D

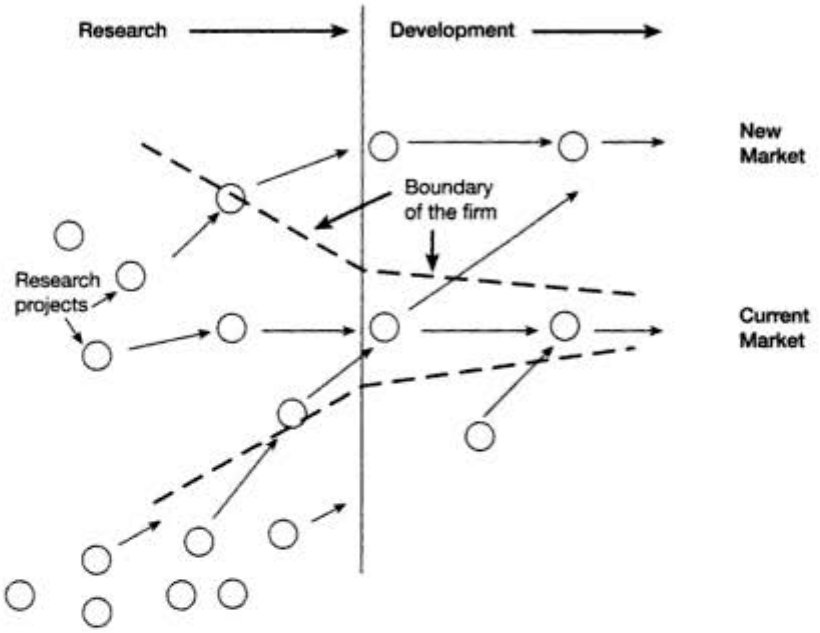

Although the Open Innovation process still weeds out false positives (now from external as well as internal sources), it also enables the recovery of false negatives, that is, projects that initially seem almost worthless, but turn out to be surprisingly valuable, as in the case of Xerox PARC noted earlier. Often these projects find value in a new market, rather than in the current market. Or they may be worthwhile if they can be combined with other projects. These opportunities were frequently overlooked by the earlier Closed Innovation process.

At root, the logic of Open Innovation is based on a landscape of abundant knowledge, which must be used readily if it is to provide value to the company that created it. The knowledge that a company uncovers in its research cannot be restricted to its internal pathways to market. Similarly, its internal pathways to market cannot necessarily be restricted to using the company's internal knowledge. This perspective suggests some very different organizing principles for research and for innovation.

Table I-1 shows some of the principles of this new paradigm and contrasts them with the earlier logic of the Closed Innovation approach.

Assessing the Prevalence of Open Innovation

The Closed Innovation paradigm has eroded in various industries. This book provides a number of detailed studies in industries such as those involving copiers, computers, disk drives, semiconductors, semiconductor equipment, communications equipment, pharmaceuticals, and biotechnology. These examples obviously all come from high-technology industries.

But don't be fooled—the concepts in this book are not specific to the high-tech portion of the overall economy. Every company has a technology, that is, a means to convert inputs into goods and services that the company sells. And no company can expect its technology to remain fixed for very long. It is far wiser to expect technology to change, sometimes in unpredictable ways, than it is to assume that things will remain in their current state for a prolonged period. Companies that don't innovate, die.

TABLE I-1

Contrasting Principles of Closed and Open Innovation

Closed Innovation Principles	Open Innovation Principles
The smart people in our field work for us.	Not all the smart people work for us. We need to work with smart people inside *and* outside our company.
To profit from R&D, we must discover it, develop it, and ship it ourselves.	External R&D can create significant value; internal R&D is needed to claim some portion of that value.
If we discover it ourselves, we will get it to market first.	We don't have to originate the research to profit from it.
The company that gets an innovation to market first will win.	Building a better business model is better than getting to market first.
If we create the most and the best ideas in the industry, we will win.	If we make the best use of internal and external ideas, we will win.
We should control our IP, so that our competitors don't profit from our ideas.	We should profit from others' use of our IP, and we should buy others' IP whenever it advances our own business model.

One example of the broad relevance of the Open Innovation approach comes from the decidedly non-high-tech consumer packaged goods industry. In 1999, Procter & Gamble decided to change its approach to innovation. The firm extended its internal R&D to the outside world through an initiative called Connect and Develop. This initiative emphasized the need for P&G to reach out to external parties for innovative ideas.[3] P&G has created a position entitled Director of External Innovation, and has set an internal goal of sourcing 50 percent of its innovations from outside the company in five years, up from an estimated 10 percent in 2002. The company's rationale is simple: Inside P&G are more than 8,600 scientists advancing the industrial knowledge that enables new P&G offerings; outside are 1.5 million. So why try to invent everything internally?[4] P&G also tries to move its own ideas outside as well. The ideas that P&G generates in its labs and that are not picked up by its internal businesses are available to other firms, even direct competitors, after three years.[5]

This is not to argue that all industries now operate in an Open Innovation regime. Some industries have not been severely impacted by the erosion factors noted previously, and they continue to operate in a Closed Innovation regime. Nuclear reactors and aircraft engines are two industries in which reliance on one's own ideas, and internal commercialization paths to market, appear to remain the dominant innovation mode. (The innovation process of designing and assembling aircraft using those engines, however, is undergoing important changes.)

Other industries have been in an Open Innovation mode for many years: The Hollywood film industry, for example, has innovated for decades through a network of partnerships and alliances between production studios, directors, talent agencies, actors, scriptwriters, specialized subcontractors (e.g., suppliers of special effects), and independent producers. Modern-day investment banking has been using external ideas for its innovations for many years as well. Newly minted Ph.D.s and even university finance professors develop new, exotic varieties of investment instruments, to hedge against risks that could not have been financed a generation ago.

These different industries can be located on a continuum, one end of which includes industries in which entirely Closed Innovation conditions prevail, the other end containing industries with fully Open Innovation conditions:[6]

Closed Innovation	*Open Innovation*
• Examples of industries: nuclear reactors, mainframe computers	• Examples of industries: PCs, movies
• Largely internal ideas	• Many external ideas
• Low labor mobility	• High labor mobility
• Little VC	• Active VC
• Few, weak start-ups	• Numerous start-ups
• Universities unimportant	• Universities important

Many industries are in transition between the two paradigms: Automotive, biotechnology, pharmaceuticals, health care, computers, software, communications, banking, insurance, consumer packaged goods, and even military weapons and communications systems are examples. It is within these transition areas that the book's concepts will be the most important. In these industries, many critically important innovations have emerged from what seemed like unlikely places. The locus of innovation in these industries is moving beyond the confines of the central R&D laboratories of the largest companies and is spreading to start-ups, to universities, and to other outsiders. If the locus of innovation is shifting in your business too, then this will be a valuable book for you.

If your industry appears to be largely buffered from the erosion factors that undermine the Closed Innovation approach, then you might expect little of value in this book. Before you close this book and place it back on the shelf, however, be sure you consider your industry carefully. Business history is full of prosperous, successful companies that were doing very well financially, even as the basis for their success was being cut out from underneath them. Some companies described in this book were also doing well for many years by adhering to the Closed Innovation model; fewer were able to detect and respond to the erosion factors they encountered before it was too late. Is it possible that your industry is also under pressure from one or more of these erosion factors and the effects just haven't materialized yet? If so, it would be wise to learn about the experience of other industries, to gain some insights about what you might do if those effects should happen to you.

Insights from the Book

When the innovation context shifts from Closed toward Open, the process of innovation must change as well. A number of insights that result from this new view of innovation will be presented in this book.

Chapter 1 presents the experience of the Xerox Corporation in dealing with a highly productive research laboratory, the Palo Alto Research Center. Xerox selected PARC technologies that fit its business model, and eschewed those that did not. These rejected technologies were later commercialized outside of Xerox's value chain, enabling instead different value chains across numerous companies. Some of the same technologies that Xerox rightly rejected as being of little value for its business model went on to become quite valuable indeed, albeit through the use of very different business models. Xerox's management of its PARC technologies illustrates in a nutshell the transition from Closed Innovation to Open Innovation.

Chapters 2 and 3 explore the Closed and Open Innovation models in more detail. This gives rise to a key insight: Useful knowledge has become widely diffused. A century ago, many leading industrial companies held knowledge monopolies; they led the industry and indeed the world in the critical discoveries that supported their industry. Today, these knowledge monopolies have been largely broken up, sometimes by government antitrust policy, but more often by the onslaught of new start-up companies, accompanied by the increasing quality and productivity of university research. The distribution of knowledge has spilled out well beyond the knowledge held by central research laboratories, with important pools of knowledge distributed among companies, customers, suppliers, universities, national labs, industry consortia, and start-up firms.

Chapter 3 also argues that companies don't take full advantage of this wealth of information. Companies often err by making too little use of others' ideas in their own businesses, causing wasteful duplication of innovative effort. This makes their internal R&D slower to achieve results, and less productive as well. Companies also often err by allowing too little use of their own ideas in others' businesses, forgoing additional profits from others' use of their ideas.

A related insight is that ideas that are not readily used can be lost. The erosion factors that undermine Closed Innovation also undermine companies' preferences to place ideas on the shelf until they can be used

internally. Such mothballing of ideas is increasingly untenable: Ideas, and the people who create them, no longer can be warehoused until the companies' own businesses are ready to make use of them. Companies that do not use their ideas with alacrity risk losing them—and the people who thought of them—to outside organizations.

Chapter 4 presents a fourth insight: The value of an idea or a technology depends on its business model. There is no inherent value in a technology per se. The value is determined instead by the business model used to bring it to market. The same technology taken to market through two different business models will yield different amounts of value. An inferior technology with a better business model will often trump a better technology commercialized through an inferior business model. The business model defines what customer problems are being solved, and looks for external and internal ideas to solve them. It also specifies how some portion of that value will be claimed. This also has implications for managing intellectual property, as we will describe later.

Chapters 5 through 8 provide detailed illustrations of Open Innovation concepts in action at leading firms. Chapter 5 recounts the transformation in how IBM manages innovation. IBM was one of the paradigmatic practitioners of Closed Innovation for most of its existence. Yet today it has shed much of the mental baggage associated with that approach, to the point that IBM frequently uses others' technologies in its business and offers its technology for sale to others to use in their business.

Chapter 6 depicts the very different innovation model of the Intel Corporation. From its very inception, Intel eschewed many ideas of the Closed Innovation paradigm. In a very high-tech industry, Intel does relatively little internal research and has organized itself instead to leverage external technologies. It does this through careful monitoring of external academic research and through corporate VC investments in external start-up companies.

In contrast to Intel's approach to bringing external technology inside, chapter 7 describes Lucent's approach to taking internal technology outside. Lucent's New Ventures Group acts as an internal VC group within Bell Labs. Its presence influences the way that Lucent commercializes Bell Laboratories technologies—both the technologies that remain within Lucent and those that go to market through a newly formed venture.

Intel and Lucent illustrate another key concept for Open Innovation: The presence of VC changes the innovation process for everyone. The impact of VC extends well beyond the start-up companies it finances.

Venture capital ultimately influences the companies that lose people to start-ups or that buy from, sell to, compete with, or partner with them. Venture capital processes for adding value to the companies they finance are not well understood in technology circles, yet these processes are critical in the Open Innovation paradigm. Established companies at a minimum must learn to coexist with VC. Ideally, they should learn to exploit VC's ability to fund multiple organizational experiments to commercialize technologies, and treat those experiments as early market explorations for their own future growth.

Chapter 8 examines the management of intellectual property (IP) in the innovation process. This chapter highlights a final insight of Open Innovation: In a world of abundant knowledge, companies should be active buyers—and active sellers—of IP. Few companies take full commercial advantage of their own IP beyond using it in their own business. And every company can benefit from utilizing external IP in its business, rather than inventing it from scratch on its own. This requires an entirely different mind-set toward managing IP: Instead of managing your IP to exclude rivals, manage your IP to profit from others' use of it. And don't be afraid to profit from others' IP in your own business. Millennium Pharmaceuticals, IBM, and Intel illustrate the fascinating opportunities that exist in managing IP.

Chapter 9 explores how companies can make the transition to a more Open Innovation system. The chapter explains how a company can exploit the principles of Open Innovation. It examines how external technologies can fill the gaps in a company's current business. It also looks at how internal technologies can generate the seeds of a company's new business.

These chapters collectively call for a new vision of the innovation process. This vision eagerly seeks external knowledge and ideas, even as it nurtures internal ones. It utilizes valuable ideas from whatever source in advancing a company's own business, and it places the company's own ideas in other companies' businesses. By opening itself up to the world of knowledge that surrounds it, the twenty-first-century corporation can avoid the innovation paradox that plagues so many firms' R&D activities today. In so doing, the company can renew its current business and generate new business. For an innovative company in a world of abundant knowledge, today can be the best of times.