

ISM105/125

The Disk Drive Industry
Examples of Functional Maps

Ref: Christensen, Clayton M.
"The Innovator's Dilemma"
(1997)

CHAPTER ONE

How Can Great Firms Fail? Insights from the Hard Disk Drive Industry



When I began my search for an answer to the puzzle of why the best firms can fail, a friend offered some sage advice. "Those who study genetics avoid studying humans," he noted. "Because new generations come along only every thirty years or so, it takes a long time to understand the cause and effect of any changes. Instead, they study fruit flies, because they are conceived, born, mature, and die all within a single day. If you want to understand why something happens in business, study the disk drive industry. Those companies are the closest things to fruit flies that the business world will ever see."

Indeed, nowhere in the history of business has there been an industry like disk drives, where changes in technology, market structure, global scope, and vertical integration have been so pervasive, rapid, and unrelenting. While this pace and complexity might be a nightmare for managers, my friend was right about its being fertile ground for research. Few industries offer researchers the same opportunities for developing theories about how different types of change cause certain types of firms to succeed or fail or for testing those theories as the industry repeats its cycles of change.

This chapter summarizes the history of the disk drive industry in all its complexity. Some readers will be interested in it for the sake of history itself.¹ But the value of understanding this history is that out of its complex-

ity emerge a few stunningly simple and consistent factors that have repeatedly determined the success and failure of the industry's best firms. Simply put, when the best firms succeeded, they did so because they listened responsively to their customers and invested aggressively in the technology, products, and manufacturing capabilities that satisfied their customers' next-generation needs. But, paradoxically, when the best firms subsequently failed, it was for the same reasons—they listened responsibly to their customers and invested aggressively in the technology, products, and manufacturing capabilities that satisfied their customers' next-generation needs. This is one of the innovator's dilemmas: Blindly following the maxim that good managers should keep close to their customers can sometimes be a fatal mistake.


The history of the disk drive industry provides a framework for understanding when "keeping close to your customers" is good advice—and when it is not. The robustness of this framework could only be explored by researching the industry's history in careful detail. Some of that detail is recounted here, and elsewhere in this book, in the hope that readers who are immersed in the detail of their own industries will be better able to recognize how similar patterns have affected their own fortunes and those of their competitors.

HOW DISK DRIVES WORK

Disk drives write and read information that computers use. They comprise read-write heads mounted at the end of an arm that swings over the surface of a rotating disk in much the same way that a phonograph needle and arm reach over a record; aluminum or glass disks coated with magnetic material; at least two electric motors, a spin motor that drives the rotation of the disks and an actuator motor that moves the head to the desired position over the disk; and a variety of electronic circuits that control the drive's operation and its interface with the computer. See Figure 1.1 for an illustration of a typical disk drive.

The read-write head is a tiny electromagnet whose polarity changes whenever the direction of the electrical current running through it changes. Because opposite magnetic poles attract, when the polarity of the head becomes positive, the polarity of the area on the disk beneath the head switches to negative, and vice versa. By rapidly changing the direction of current flowing through the head's electromagnet as the disk spins beneath

Figure 1.1 Primary Components of a Typical Disk Drive




the head, a sequence of positively and negatively oriented magnetic domains are created in concentric tracks on the disk's surface. Disk drives can use the positive and negative domains on the disk as a binary numeric system—1 and 0—to "write" information onto disks. Drives read information from disks in essentially the opposite process: Changes in the magnetic flux fields on the disk surface induce changes in the micro current flowing through the head.

EMERGENCE OF THE EARLIEST DISK DRIVES

A team of researchers at IBM's San Jose research laboratories developed the first disk drive between 1952 and 1956. Named RAMAC (for Random Access Method for Accounting and Control), this drive was the size of a large refrigerator, incorporated fifty twenty-four-inch disks, and could store 5 megabytes (MB) of information (see Figure 1.2). Most of the fundamental architectural concepts and component technologies that defined today's dominant disk drive design were also developed at IBM. These include its removable packs of rigid disks (introduced in 1961);

Figure 1.2 The First Disk Drive, Developed by IBM



Source: Courtesy of International Business Machines Corporation.

the floppy disk drive (1971); and the Winchester architecture (1973). All had a powerful, defining influence on the way engineers in the rest of the industry defined what disk drives were and what they could do.

As IBM produced drives to meet its own needs, an independent disk drive industry emerged serving two distinct markets. A few firms developed the plug-compatible market (PCM) in the 1960s, selling souped-up copies of IBM drives directly to IBM customers at discount prices. Although most of IBM's competitors in computers (for example, Control Data, Burroughs, and Univac) were integrated vertically into the manufacture of their own disk drives, the emergence in the 1970s of smaller, nonintegrated computer makers such as Nixdorf, Wang, and Prime spawned an original equipment market (OEM) for disk drives as well. By 1976 about \$1 billion worth of disk drives were produced, of which

captive production accounted for 50 percent and PCM and OEM for about 25 percent each.

The next dozen years unfolded a remarkable story of rapid growth, market turbulence, and technology-driven performance improvements. The value of drives produced rose to about \$18 billion by 1995. By the mid-1980s the PCM market had become insignificant, while OEM output grew to represent about three-fourths of world production. Of the seven-teen firms populating the industry in 1976—all of which were relatively large, diversified corporations such as Diablo, Ampex, Memorex, EMM, and Control Data—all except IBM's disk drive operation had failed or had been acquired by 1995. During this period an additional 129 firms entered the industry, and 109 of those also failed. Aside from IBM, Fujitsu, Hitachi, and NEC, all of the producers remaining by 1996 had entered the industry as start-ups after 1976.


Some have attributed the high mortality rate among the integrated firms that created the industry to its nearly unfathomable pace of technological change. Indeed, the pace of change has been breathtaking. The number of megabits (Mb) of information that the industry's engineers have been able to pack into a square inch of disk surface has increased by 35 percent per year, on average, from 50 Kb in 1967 to 1.7 Mb in 1973, 12 Mb in 1981, and 1100 Mb by 1995. The physical size of the drives was reduced at a similar pace: The smallest available 20 MB drive shrank from 800 cubic inches (in.³) in 1978 to 1.4 in.³ by 1993—a 35 percent annual rate of reduction.

Figure 1.3 shows that the slope of the industry's experience curve (which correlates the cumulative number of terabytes (one thousand gigabytes) of disk storage capacity shipped in the industry's history to the constant-dollar price per megabyte of memory) was 53 percent—meaning that with each doubling of cumulative terabytes shipped, cost per megabyte fell to 53 percent of its former level. This is a much steeper rate of price decline than the 70 percent slope observed in the markets for most other microelectronics products. The price per megabyte has declined at about 5 percent per *quarter* for more than twenty years.

THE IMPACT OF TECHNOLOGICAL CHANGE

My investigation into why leading firms found it so difficult to stay atop the disk drive industry led me to develop the “technology mudslide hypothesis”: Coping with the relentless onslaught of technology change

Figure 1.3 Disk Drive Price Experience Curve



Source: Data are from various issues of *Disk/Trend Report*.

was akin to trying to climb a mudslide raging down a hill. You have to scramble with everything you've got to stay on top of it, and if you ever once stop to catch your breath, you get buried.

To test this hypothesis, I assembled and analyzed a database consisting of the technical and performance specifications of every model of disk drive introduced by every company in the world disk drive industry for each of the years between 1975 and 1994.² This database enabled me to identify the firms that led in introducing each new technology; to trace how new technologies were diffused through the industry over time; to see which firms led and which lagged; and to measure the impact each technological innovation had on capacity, speed, and other parameters of disk drive performance. By carefully reconstructing the history of each technological change in the industry, the changes that catapulted entrants to success or that precipitated the failure of established leaders could be identified.

This study led me to a very different view of technology change than the work of prior scholars on this question had led me to expect. Essentially, it revealed that neither the pace nor the difficulty of technological change

lay at the root of the leading firms' failures. The technology mudslide hypothesis was wrong.

The manufacturers of most products have established a trajectory of performance improvement over time.³ Intel, for example, pushed the speed of its microprocessors ahead by about 20 percent per year, from its 8 megahertz (MHz) 8088 processor in 1979 to its 133 MHz Pentium chip in 1994. Eli Lilly and Company improved the purity of its insulin from 50,000 impure parts per million (ppm) in 1925 to 10 ppm in 1980, a 14 percent annual rate of improvement. When a measurable trajectory of improvement has been established, determining whether a new technology is likely to improve a product's performance relative to earlier products is an unambiguous question.

But in other cases, the impact of technological change is quite different. For instance, is a notebook computer better than a mainframe? This is an ambiguous question because the notebook computer established a completely new performance trajectory, with a definition of performance that differs substantially from the way mainframe performance is measured. Notebooks, as a consequence, are generally sold for very different uses.

This study of technological change over the history of the disk drive industry revealed two types of technology change, each with very different effects on the industry's leaders. Technologies of the first sort *sustained* the industry's rate of improvement in product performance (total capacity and recording density were the two most common measures) and ranged in difficulty from incremental to radical. The industry's dominant firms always led in developing and adopting these technologies. By contrast, innovations of the second sort *disrupted* or redefined performance trajectories—and consistently resulted in the failure of the industry's leading firms.⁴

The remainder of this chapter illustrates the distinction between sustaining and disruptive technologies by describing prominent examples of each and summarizing the role these played in the industry's development. This discussion focuses on differences in how established firms came to lead or lag in developing and adopting new technologies, compared with entrant firms. To arrive at these examples, each new technology in the industry was examined. In analyzing which firms led and lagged at each of these points of change, I defined *established firms* to be those that had been established in the industry prior to the advent of the technology in question, practicing the prior technology. I defined *entrant firms* as those


that were new to the industry at that point of technology change. Hence, a given firm would be considered an entrant at one specific point in the industry's history, for example, at the emergence of the 8-inch drive. Yet the same firm would be considered an established firm when technologies that emerged subsequent to the firm's entry were studied.

SUSTAINING TECHNOLOGICAL CHANGES

In the history of the disk drive industry, most technology changes have sustained or reinforced established trajectories of product performance improvement. Figure 1.4, which compares the average recording density of drives that employed successive generations of head and disk technologies, maps an example of this. The first curve plots the density of drives that used conventional particulate oxide disk technology and ferrite head technology; the second charts the average density of drives that used new-technology thin-film heads and disks; the third marks the improvements in density achievable with the latest head technology, magneto-resistive heads.⁵

The way such new technologies as these emerge to surpass the perform-

Figure 1.4 Impact of New Read-Write Head Technologies in Sustaining the Trajectory of Improvement in Recording Density



Source: Data are from various issues of *Disk/Trend Report*.


ance of the old resembles a series of intersecting technology S-curves.⁶ Movement along a given S-curve is generally the result of incremental improvements within an existing technological approach, whereas jumping onto the next technology curve implies adopting a radically new technology. In the cases measured in Figure 1.4, incremental advances, such as grinding the ferrite heads to finer, more precise dimensions and using smaller and more finely dispersed oxide particles on the disk's surface, led to the improvements in density from 1 to 20 megabits per square inch (Mbps) between 1976 and 1989. As S-curve theory would predict, the improvement in recording density obtainable with ferrite/oxide technology began to level off toward the end of the period, suggesting a maturing technology. The thin-film head and disk technologies' effect on the industry sustained performance improvement at its historical rate. Thin-film heads were barely established in the early 1990s, when even more advanced magneto-resistive head technology emerged. The impact of magneto-resistive technology sustained, or even accelerated, the rate of performance improvement.

Figure 1.5 describes a sustaining technological change of a very different character: an innovation in product architecture, in which the 14-inch Winchester drive is substituted for removable disk packs, which had been the dominant design between 1962 and 1978. Just as in the thin-film for ferrite/oxide substitution, the impact of Winchester technology sustained the historically established rate of performance improvement. Similar graphs could be constructed for most other technological innovations in the industry, such as embedded servo systems, RLL and PRML recording codes, higher RPM motors, and embedded interfaces. Some of these were straightforward technology improvements; others were radical departures. But all had a similar impact on the industry: They helped manufacturers to sustain the rate of historical performance improvement that their customers had come to expect.⁷

In literally every case of sustaining technology change in the disk drive industry, established firms led in development and commercialization. The emergence of new disk and head technologies illustrates this.

In the 1970s, some manufacturers sensed that they were reaching the limit on the number of bits of information they could pack onto oxide disks. In response, disk drive manufacturers began studying ways of applying super-thin films of magnetic metal on aluminum to sustain the historical rate of improvements in recording density. The use of thin-film coatings was then highly developed in the integrated circuit industry, but

Figure 1.5 Sustaining Impact of the Winchester Architecture on the Recording Density of 14-inch Disk Drives



Source: Data are from various issues of *Disk/Trend Report*.

its application to magnetic disks still presented substantial challenges. Experts estimate that the pioneers of thin-film disk technology—IBM, Control Data, Digital Equipment, Storage Technology, and Ampex—each took more than eight years and spent more than \$50 million in that effort. Between 1984 and 1986, about two-thirds of the producers active in 1984 introduced drives with thin-film disks. The overwhelming majority of these were established industry incumbents. Only a few entrant firms attempted to use thin-film disks in their initial products, and most of those folded shortly after entry.

The same pattern was apparent in the emergence of thin-film heads. Manufacturers of ferrite heads saw as early as 1965 the approaching limit to improvements in this technology; by 1981 many believed that the limits of precision would soon be reached. Researchers turned to thin-film technology, produced by sputtering thin films of metal on the recording head and then using photolithography to etch much finer elec-

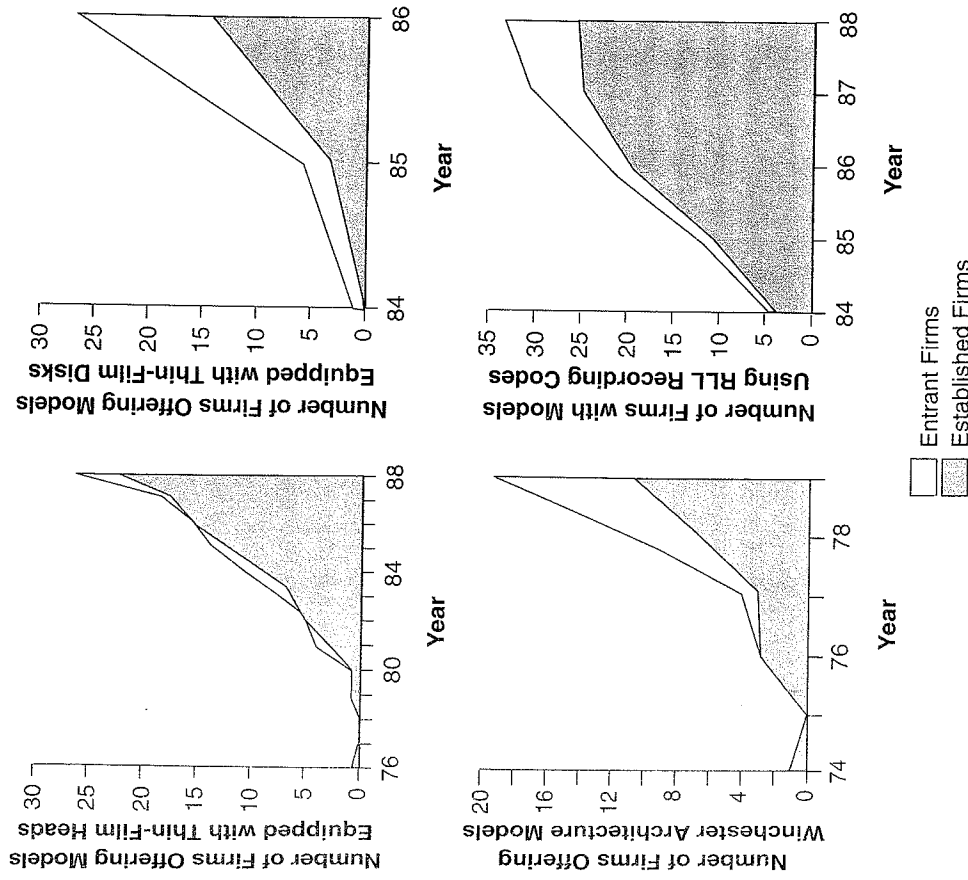
tromagnets than could be attained with ferrite technology. Again, this proved extraordinarily difficult. Burroughs in 1976, IBM in 1979, and other established firms first successfully incorporated thin-film heads in disk drives. In the period between 1982 and 1986, during which some sixty firms entered the rigid disk drive industry, only four (all commercial failures) attempted to do so using thin-film heads in their initial products as a source of performance advantage. All other entrant firms—even aggressively performance-oriented firms such as Maxtor and Conner Peripherals—found it preferable to learn their way using conventional ferrite heads first, before tackling thin-film technology.

As was the case with thin-film disks, the introduction of thin-film heads entailed the sort of sustained investment that only established firms could handle. IBM and its rivals each spent more than \$100 million developing thin-film heads. The pattern was repeated in the next-generation magnetoresistive head technology: The industry's largest firms—IBM, Seagate, and Quantum—led the race.

The established firms were the leading innovators not just in developing risky, complex, and expensive component technologies such as thin-film heads and disks, but in *literally every other one of the sustaining innovations in the industry's history*. Even in relatively simple innovations, such as RLL recording codes (which took the industry from double- to triple-density disks), established firms were the successful pioneers, and entrant firms were the technology followers. This was also true for those architectural innovations—for example, 14-inch and 2.5-inch Winchester drives—whose impact was to sustain established improvement trajectories. Established firms beat out the entrants.

Figure 1.6 summarizes this pattern of technology leadership among established and entrant firms offering products based on new sustaining technologies during the years when those technologies were emerging. The pattern is stunningly consistent. Whether the technology was radical or incremental, expensive or cheap, software or hardware, component or architecture, competence-enhancing or competence-destroying, the pattern was the same. When faced with sustaining technology change that gave existing customers something more and better in what they wanted, the leading practitioners of the prior technology led the industry in the development and adoption of the new. Clearly, the leaders in this industry did not fail because they became passive, arrogant, or risk-averse or because they couldn't keep up with the stunning rate of technological change. My technology mudslide hypothesis wasn't correct.

Figure 1.6 Leadership of Established Firms in Sustaining Technologies



Source: Data are from various issues of *Disk/Trend Report*.

FAILURE IN THE FACE OF DISRUPTIVE TECHNOLOGICAL CHANGES

Most technological change in the disk drive industry has consisted of sustaining innovations of the sort described above. In contrast, there have been only a few of the other sort of technological change, called disruptive technologies. These were the changes that toppled the industry's leaders.

The most important disruptive technologies were the architectural innovations that shrunk the size of the drives—from 14-inch diameter disks to diameters of 8, 5.25, and 3.5-inches and then from 2.5 to 1.8 inches. Table 1.1 illustrates the ways these innovations were disruptive. Based on 1981 data, it compares the attributes of a typical 5.25-inch drive, a new architecture that had been in the market for less than a year, with those of a typical 8-inch drive, which at that time was the standard drive used by minicomputer manufacturers. Along the dimensions of performance important to established minicomputer manufacturers—capacity, cost per megabyte, and access time—the 8-inch product was vastly superior. The 5.25-inch architecture did not address the perceived needs of minicomputer manufacturers at that time. On the other hand, the 5.25-inch drive had features that appealed to the desktop personal computer market segment just emerging in the period between 1980 and 1982. It was small and lightweight, and, priced at around \$2,000, it could be incorporated into desktop machines economically.

Generally disruptive innovations were technologically straightforward, consisting of off-the-shelf components put together in a product architecture that was often simpler than prior approaches.⁸ They offered less of what customers in established markets wanted and so could rarely be initially employed there. They offered a different package of attributes valued only in emerging markets remote from, and unimportant to, the mainstream.


The trajectory map in Figure 1.7 shows how this series of simple but disruptive technologies proved to be the undoing of some very aggressive,

Table 1.1 A Disruptive Technology Change: The 5.25-inch Winchester Disk Drive (1981)

Attribute	8-Inch Drives (Minicomputer Market)	5.25-Inch Drives (Desktop Computer Market)
Capacity (megabytes)	60	10
Physical volume (cubic inches)	566	150
Weight (pounds)	21	6
Access time (milliseconds)	30	160
Cost per megabyte	\$50	\$200
Unit cost	\$3000	\$2000

Source: Data are from various issues of *Disk/Trend Report*.

Figure 1.7 Intersecting Trajectories of Capacity Demanded versus Capacity Supplied in Rigid Disk Drives



Source: Clayton M. Christensen, "The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence," *Business History Review* 67, no. 4 (Winter 1993): 559. Reprinted by permission.

astutely managed disk drive companies. Until the mid-1970s, 14-inch drives with removable packs of disks accounted for nearly all disk drive sales. The 14-inch Winchester architecture then emerged to sustain the trajectory of recording density improvement. Nearly all of these drives (removable disks and Winchesters) were sold to mainframe computer

manufacturers, and the same companies that led the market in disk pack drives led the industry's transition to the Winchester technology.

The trajectory map shows that the hard disk capacity provided in the median priced, typically configured mainframe computer system in 1974 was about 130 MB per computer. This increased at a 15 percent annual rate over the next fifteen years—a trajectory representing the disk capacity demanded by the typical users of new mainframe computers. At the same time, the capacity of the average 14-inch drive introduced for sale each year increased at a faster, 22 percent rate, reaching beyond the mainframe market to the large scientific and supercomputer markets.⁹

Between 1978 and 1980, several entrant firms—Shugart Associates, Micropolis, Priam, and Quantum—developed smaller 8-inch drives with 10, 20, 30, and 40 MB capacity. These drives were of no interest to mainframe computer manufacturers, which at that time were demanding drives with 300 to 400 MB capacity. These 8-inch entrants therefore sold their disruptive drives into a new application—minicomputers.¹⁰ The customers—Wang, DEC, Data General, Prime, and Hewlett-Packard—did not manufacture mainframes, and their customers often used software substantially different from that used in mainframes. These firms hitherto had been unable to offer disk drives in their small, desk-side minicomputers because 14-inch models were too big and expensive. Although initially the cost per megabyte of capacity of 8-inch drives was higher than that of 14-inch drives, these new customers were willing to pay a premium for other attributes that were important to them—especially smaller size. Smallness had little value to mainframe users.

Once the use of 8-inch drives became established in minicomputers the hard disk capacity shipped with the median-priced minicomputer grew about 25 percent per year: a trajectory determined by the ways in which minicomputer owners learned to use their machines. At the same time however, the 8-inch drive makers found that, by aggressively adopting sustaining innovations, they could increase the capacity of their product at a rate of more than 40 percent per year—nearly double the rate of increase demanded by their original "home" minicomputer market. In consequence, by the mid-1980s, 8-inch drive makers were able to provide the capacities required for lower-end mainframe computers. Unit volume had grown significantly so that the cost per megabyte of 8-inch drive had declined below that of 14-inch drives, and other advantages became apparent: For example, the same percentage mechanical vibration in an 8-inch drive, as opposed to a 14-inch drive, caused much less variance in

the absolute position of the head over the disk. Within a three-to-four-year period, therefore, 8-inch drives began to invade the market above them, substituting for 14-inch drives in the lower-end mainframe computer market.

As the 8-inch products penetrated the mainframe market, the established manufacturers of 14-inch drives began to fail. Two-thirds of them never introduced an 8-inch model. The one-third that introduced 8-inch models did so about two years behind the 8-inch entrant manufacturers. Ultimately, every 14-inch drive maker was driven from the industry.¹¹

The 14-inch drive makers were not toppled by the 8-inch entrants because of technology. The 8-inch products generally incorporated standard off-the-shelf components, and when those 14-inch drive makers that did introduce 8-inch models got around to doing so, their products were very performance-competitive in capacity, areal density, access time, and price per megabyte. The 8-inch models introduced by the established firms in 1981 were nearly identical in performance to the average of those introduced that year by the entrant firms. In addition, the rates of improvement in key attributes (measured between 1979 and 1983) were stunningly similar between established and entrant firms.¹²

Held Captive by Their Customers

Why were the leading drive makers unable to launch 8-inch drives until it was too late? Clearly, they were technologically capable of producing these drives. Their failure resulted from delay in making the strategic commitment to enter the emerging market in which the 8-inch drives initially could be sold. Interviews with marketing and engineering executives close to these companies suggest that the established 14-inch drive manufacturers were held captive by customers. Mainframe computer manufacturers did not need an 8-inch drive. In fact, they explicitly did not want it: they wanted drives with increased capacity at a lower cost per megabyte. The 14-inch drive manufacturers were listening and responding to their established customers. And their customers—in a way that was not apparent to either the disk drive manufacturers or their computer-making customers—were pulling them along a trajectory of 22 percent capacity growth in a 14-inch platform that would ultimately prove fatal.¹³

Figure 1.7 maps the disparate trajectories of performance improvement demanded in the computer product segments that emerged later, compared to the capacity that changes in component technology and refinements

in system design made available within each successive architecture. The solid lines emanating from points A, B, C, D, and E measure the disk drive capacity provided with the median-priced computer in each category, while the dotted lines from the same points measure the average capacity of all disk drives introduced for sale in each architecture, for each year. These transitions are briefly described below.

The Advent of the 5.25-inch Drive

In 1980, Seagate Technology introduced 5.25-inch disk drives. Their capacities of 5 and 10 MB were of no interest to minicomputer manufacturers, who were demanding drives of 40 and 60 MB from their suppliers. Seagate and other firms that entered with 5.25-inch drives in the period 1980 to 1983 (for example, Miniscribe, Computer Memories, and International Memories) had to pioneer new applications for their products and turned primarily to desktop personal computer makers. By 1990, the use of hard drives in desktop computers was an obvious application for magnetic recording. It was not at all clear in 1980, however, when the market was just emerging, that many people could ever afford or use a hard drive on the desktop. The early 5.25-inch drive makers found this application (one might even say that they *enabled* it) by trial and error, selling drives to whomever would buy them.

Once the use of hard drives was established in desktop PCs, the disk capacity shipped with the median-priced machine (that is, the capacity demanded by the general PC user) increased about 25 percent per year. Again, the technology improved at nearly twice the rate demanded in the new market: The capacity of new 5.25-inch drives increased about 50 percent per year between 1980 and 1990. As in the 8-inch for 14-inch substitution, the first firms to produce 5.25-inch drives were entrants; on average, established firms lagged behind entrants by two years. By 1985, only half of the firms producing 8-inch drives had introduced 5.25-inch models. The other half never did.

Growth in the use of 5.25-inch drives occurred in two waves. The first followed creation of a new application for rigid disk drives: desktop computing, in which product attributes such as physical size, relatively unimportant in established applications, were highly valued. The second wave followed substitution of 5.25-inch disks for larger drives in established minicomputer and mainframe computer markets, as the rapidly increasing capacity of 5.25-inch drives intersected the more slowly grow-

ing trajectories of capacity demanded in these markets. Of the four leading 8-inch drive makers—Shugart Associates, Micropolis, Priam, and Quantum—only Micropolis survived to become a significant manufacturer of 5.25-inch drives, and that was accomplished only with Herculean managerial effort, as described in chapter 5.

The Pattern Is Repeated: The Emergence of the 3.5-inch Drive

The 3.5-inch drive was first developed in 1984 by Rodime, a Scottish entrant. Sales of this architecture were not significant, however, until Conner Peripherals, a spinoff of 5.25-inch drive makers Seagate and Miniscribe, started shipping product in 1987. Conner had developed a small, lightweight drive architecture that was much more rugged than its 5.25-inch ancestors. It handled electronically functions that had previously been managed with mechanical parts, and it used microcode to replace functions that had previously been addressed electronically. Nearly all of Conner's first year revenues of \$113 million¹⁴ came from Compaq Computer, which had aided Conner's start-up with a \$30 million investment. The Conner drives were used primarily in a new application—portable and laptop machines, in addition to “small footprint” desktop models—where customers were willing to accept lower capacities and higher costs per megabyte to get lighter weight, greater ruggedness, and lower power consumption.

Seagate engineers were not oblivious to the coming of the 3.5-inch architecture. Indeed, in early 1985, less than one year after Rodime introduced the first 3.5-inch drive and two years before Conner Peripherals started shipping its product, Seagate personnel showed working 3.5-inch prototype drives to customers for evaluation. The initiative for the new program came from Seagate's engineering organization. Opposition to the executive team; they argued that the market wanted higher capacity drives at a lower cost per megabyte and that 3.5-inch drives could never be built at a lower cost per megabyte than 5.25-inch drives.

Seagate's marketers tested the 3.5-inch prototypes with customers in the desktop computing market it already served—manufacturers like IBM, and value-added resellers of full-sized desktop computer systems. Not surprisingly, they indicated little interest in the smaller drive. They were looking for capacities of 40 and 60 megabytes for their next-generation machines, while the 3.5-inch architecture could provide only 20 MB—and at higher costs.¹⁵

In response to lukewarm reviews from customers, Seagate's program manager lowered his 3.5-inch sales estimates, and the firm's executives canceled the program. Their reasoning? The markets for 5.25-inch products were larger, and the sales generated by spending the engineering effort on new 5.25-inch products would create greater revenues for the company than would efforts targeted at new 3.5-inch products.

In retrospect, it appears that Seagate executives read the market—at least their own market—very accurately. With established applications and product architectures of their own, such as the IBM XT and AT, these customers saw no value in the improved ruggedness or the reduced size, weight, and power consumption of 3.5-inch products.

Seagate finally began shipping 3.5-inch drives in early 1988—the same year in which the performance trajectory of 3.5-inch drives (shown in Figure 1.7) intersected the trajectory of capacity demanded in desktop computers. By that time, the industry had shipped, cumulatively, nearly \$750 million in 3.5-inch products. Interestingly, according to industry observers, as of 1991 almost none of Seagate's 3.5-inch products had been sold to manufacturers of portable/laptop/notebook computers. In other words, Seagate's primary customers were still desktop computer manufacturers, and many of its 3.5-inch drives were shipped with frames for mounting them in computers designed for 5.25-inch drives.

The fear of cannibalizing sales of existing products is often cited as a reason why established firms delay the introduction of new technologies. As the Seagate-Conner experience illustrates, however, if new technologies enable new market applications to emerge, the introduction of new technology may not be inherently cannibalistic. But when established firms wait until a new technology has become commercially mature in its new applications and launch their own version of the technology only in response to an attack on their home markets, the fear of cannibalization can become a self-fulfilling prophecy.

Although we have been looking at Seagate's response to the development of the 3.5-inch drive architecture, its behavior was not atypical; by 1988, only 35 percent of the drive manufacturers that had established themselves making 5.25-inch products for the desktop PC market had introduced 3.5-inch drives. Similar to earlier product architecture transitions, the barrier to development of a competitive 3.5-inch product does not appear to have been engineering-based. As in the 14- to 8-inch transition, the new-architecture drives introduced by the incumbent, established firms during the transitions from 8 to 5.25 inches and from 5.25 to 3.5 inches were fully performance-competitive with those of entrant drives. Rather,

the 5.25-inch drive manufacturers seem to have been misled by their customers, notably IBM and its direct competitors and resellers, who themselves seemed as oblivious as Seagate to the potential benefits and possibilities of portable computing and the new disk drive architecture that might facilitate it.

PrairieTek, Conner, and the 2.5-inch Drive

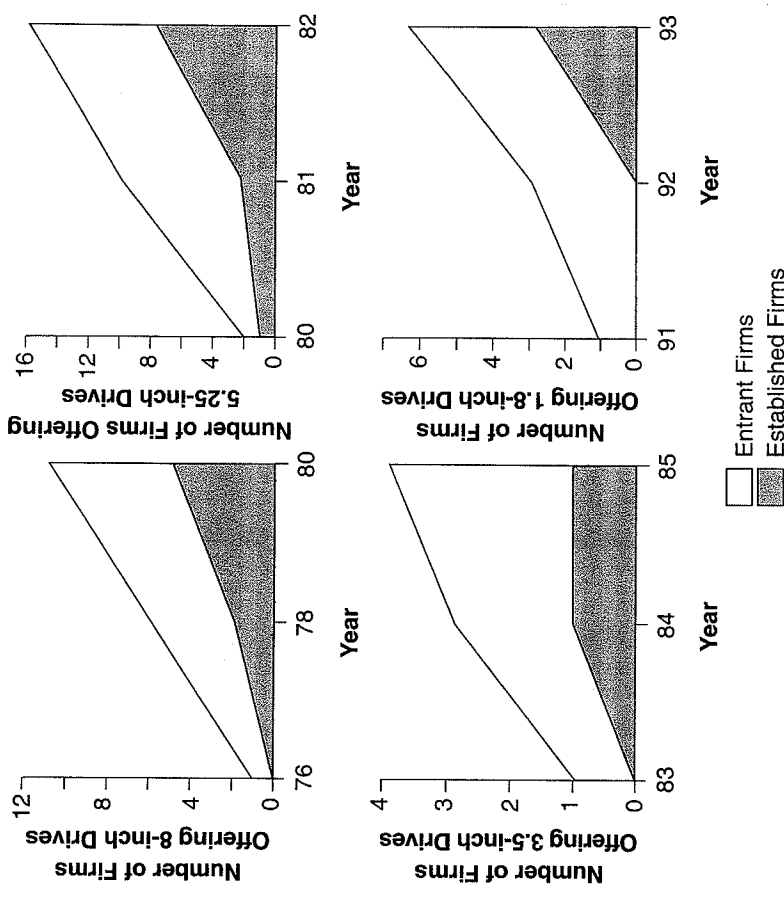
In 1989 an industry entrant in Longmont, Colorado, PrairieTek, upstaged the industry by announcing a 2.5-inch drive, capturing nearly all \$30 million of this nascent market. But Conner Peripherals announced its own 2.5-inch product in early 1990 and by the end of that year had claimed 95 percent of the 2.5-inch drive market. PrairieTek declared bankruptcy in late 1991, by which time each of the other 3.5-inch drivemakers—Quantum, Seagate, Western Digital, and Maxtor—had introduced 2.5-inch drives of their own.

What had changed? Had the incumbent leading firms finally learned the lessons of history? Not really. Although Figure 1.7 shows the 2.5-inch drive had significantly less capacity than the 3.5-inch drives, the portable computing markets into which the smaller drives were sold valued other attributes: weight, ruggedness, low power consumption, small physical size, and so on. Along these dimensions, the 2.5-inch drive offered improved performance over that of the 3.5-inch product: It was a *smart* technology. In fact, the computer makers who bought Conner's 3.5-inch drive—laptop computer manufacturers such as Toshiba, Zenith, and Sharp—were the leading makers of notebook computers, and these firms needed the smaller 2.5-inch drive architecture. Hence, Conner and its competitors in the 3.5-inch market followed their customers seamlessly across the transition to 2.5-inch drives.

In 1992, however, the 1.8-inch drive emerged, with a distinctly disruptive character. Although its story will be recounted in detail later, it suffices to state here that by 1995, it was *entrant* firms that controlled 98 percent of the \$130 million 1.8-inch drive market. Moreover, the largest initial market for 1.8-inch drives wasn't in computing at all. It was in portable heart monitoring devices!

Figure 1.8 summarizes this pattern of entrant firms' leadership in disruptive technology. It shows, for example, that two years after the 8-inch drive was introduced, two-thirds of the firms producing it (four of six), were entrants. And, two years after the first 5.25-inch drive was intro-

Figure 1.8 Leadership of Entrant Firms in Disruptive Technology



Source: Data are from various issues of *Disk/Trend Report*.

duced, 80 percent of the firms producing these disruptive drives were entrants.

SUMMARY

There are several patterns in the history of innovation in the disk drive industry. The first is that the disruptive innovations were technologically straightforward. They generally packaged known technologies in a unique architecture and enabled the use of these products in applications where magnetic data storage and retrieval previously had not been technologically or economically feasible.

The second pattern is that the purpose of advanced technology develop-

ment in the industry was always to *sustain* established trajectories of performance improvement: to reach the higher-performance, higher-margin domain of the upper right of the trajectory map. Many of these technologies were radically new and difficult, but they were not disruptive. The customers of the leading disk drive suppliers led them toward these achievements. Sustaining technologies, as a result, did not precipitate failure.

The third pattern shows that, despite the established firms' technological prowess in leading sustaining innovations, from the simplest to the most radical, the firms that led the industry in every instance of developing and adopting disruptive technologies were entrants to the industry, not its incumbent leaders.

This book began by posing a puzzle: Why was it that firms that could be esteemed as aggressive, innovative, customer-sensitive organizations could ignore or attend belatedly to technological innovations with enormous strategic importance? In the context of the preceding analysis of the disk drive industry, this question can be sharpened considerably. The established firms were, in fact, aggressive, innovative, and customer-sensitive in their approaches to sustaining innovations of every sort. But the problem established firms seem unable to confront successfully is that of *downward* vision and mobility, in terms of the trajectory map. Finding new applications and markets for these new products seems to be a capability that each of these firms exhibited once, upon entry, and then apparently lost. It was as if the leading firms were held captive by their customers, enabling attacking entrant firms to topple the incumbent industry leaders each time a disruptive technology emerged.¹⁶ Why this happened, and is still happening, is the subject of the next chapter.

APPENDIX 1.1: A NOTE ON THE DATA AND METHOD USED TO GENERATE FIGURE 1.7

The trajectories mapped in Figure 1.7 were calculated as follows. Data on the capacity provided with computers was obtained from *Data Sources*, an annual publication listing the technical specifications of all computer models available from every computer manufacturer. For instances in which particular models were available with different features and configurations, the manufacturer provided *Data Sources* with a "typical" system configuration with defined random access memory (RAM) capacity, performance specifications of peripheral equip-

ment (including disk drives), list prices, and year of introduction. For instances in which a given computer model was offered for sale over a sequence of years, the hard disk capacity provided in the typical configuration typically increased. *Data Sources* used the categories mainframe, mini/midrange, desktop personal, portable and laptop, and notebook. As of 1993, 1.8-inch drives were not being used in hand-held computers, so no data on that potential market existed.

For Figure 1.7, for each year and each class of computers, all models available for sale were ranked by price and the hard disk capacity provided with the median-priced model identified. The best-fit lines through the resultant time series were plotted as the solid lines in Figure 1.7 for expository simplification to indicate the trend in typical machines. In reality, of course, there is a wide band around these lines. The *frontier* of performance—the highest capacity offered with the most expensive computers—was substantially higher than the typical values shown.

The dotted lines in Figure 1.7 represent the best-fit line through the unweighted average capacity of all disk drives introduced for sale in each given architecture for each year. This data was taken from *Disk/Trend Report*. Again, for expository simplification, only this average line is shown. There was a wide band of capacities introduced for sale in each year, so that the frontier or highest capacity drive introduced in each year was substantially above the average shown. Stated in another way, a distinction must be made between the full range of products available for purchase and those in typical systems. The upper and lower bands around the median and average figures shown in Figure 1.7 are generally parallel to the lines shown.

Because higher capacity drives were available in the market than were offered with the median-priced systems, the solid-line trajectories in Figure 1.7, as I state in the text, represent the capacities "demanded" in each market. In other words, the capacity per machine was not constrained by technological availability. Rather, it represents the selection of hard disk capacity by computer users, given the prevailing cost.

NOTES

1. A more complete history of the disk drive industry can be found in Clayton M. Christensen, "The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence," *Business History Review* (67), Winter, 1993, 531–588. This history focuses only on the manufacturers of rigid disk drives or hard drives—products on which data are stored on rigid metal platters. Companies manufacturing floppy disk drives (removable diskettes of flexible mylar coated with iron oxide on which data are stored) historically were different firms from those making hard disk drives.

2. Much of the data for this analysis came from *Disk/Trend Report*, a highly respected annual market research publication, augmented with more detailed product-specification sheets obtained from the disk drive manufacturers themselves. I am grateful to the editors and staff at Disk/Trend, Inc., for their patient and generous assistance in this project.
3. The concept of trajectories of technological progress was examined by Giovanni Dosi in "Technological Paradigms and Technological Trajectories," *Research Policy* (11), 1982, 147-162.
4. The ways in which the findings of this study differ from those of some earlier scholars of technology change while building upon those of others are discussed in greater detail in chapter 2.
5. The first technology for making heads built an electromagnet by wrapping a fine thread of copper wire around a core of iron oxide (ferrite); hence the term *ferrite head*. Incremental improvements to this approach involved learning to grind the ferrite to finer and finer dimensions, using better lapping techniques, and strengthening the ferrite by doping it with barium. *Thin-film heads* were made photolithographically, using technology similar to that used in making integrated circuits on silicon wafers to etch the electromagnet on the surface of the head. This was difficult because it involved much thicker layers of material than were common in IC manufacturing. The third technology, adopted starting in the mid-1990s, was called *magneto-resistive heads*. These were also made with thin-film photolithography, but used the principle that changes in the magnetic flux field on the disk surface changed the electrical resistivity of the circuitry in the head. By measuring changes in resistivity rather than changes in the direction of current flow, magnetoresistive heads were much more sensitive, and hence permitted denser data recording, than prior technology. In the evolution of disk technology, the earliest disks were made by coating fine needle-shaped particles of iron oxide—literally rust—over the surface of a flat, polished aluminum platter. Hence, these disks were called *oxide* disks. Incremental improvements to this technology involved making finer and finer iron oxide particles, and dispersing them more uniformly, with fewer uncoated voids on the aluminum platter's surface. This was supplanted by a sputtering technology, also borrowed from semiconductor processing, that coated the aluminum platter with a thin film of metal a few angstroms thick. The thinness of this layer; its continuous, rather than particulate nature; and the process's flexibility in depositing magnetic materials with higher coercivity, enabled denser recording on thin-film disks than was feasible on oxide disks.
6. Richard J. Foster, *Innovation: The Attacker's Advantage* (New York: Summit Books, 1986).
7. The examples of technology change presented in Figures 1.1 and 1.2 introduce

some ambiguity to the unqualified term *discontinuity*, as used by Giovanni Dosi (see "Technological Paradigms and Technological Trajectories," *Research Policy* [11] 1982), Michael L. Tushman and Philip Anderson (see "Technological Discontinuities and Organizational Environments," *Administrative Science Quarterly* [31], 1986), and others. The innovations in head and disk technology described in Figure 1.4 represent *positive* discontinuities in an established technological trajectory, while the trajectory-disrupting technologies charted in Figure 1.7 represent *negative* discontinuities. As will be shown below, established firms seemed quite capable of leading the industry over positive discontinuities, but generally lost their industry lead when faced with negative discontinuities.

8. This tendency consistently appears across a range of industries. Richard S. Rosenbloom and Clayton M. Christensen (in "Technological Discontinuities, Organizational Capabilities, and Strategic Commitments," *Industrial and Corporate Change* [3], 1994, 655-685) suggest a much broader set of industries in which leading firms may have been toppled by technologically straightforward disruptive innovations than is covered in this book.
9. A summary of the data and procedures used to generate Figure 1.7 is included in Appendix 1.1.
10. The microcomputer market was not new in 1978, but it was a new application for Winchester-technology disk drives.
11. This statement applies only to independent drive makers competing in the OEM market. Some of the vertically integrated computer makers, such as IBM, have survived across these generations with the benefit of a captive internal market. Even IBM, however, addressed the sequence of different emerging markets for disk drives by creating autonomous "start-up" disk drive organizations to address each one. Its San Jose organization focused on high-end (primarily mainframe) applications. A separate division in Rochester, MN, focused on mid-range computers and workstations. IBM created a different organization in Fujisawa, Japan, to produce drives for the desktop personal computer market.
12. This result is very different from that observed by Rebecca M. Henderson (see *The Failure of Established Firms in the Face of Technological Change: A Study of the Semiconductor Photolithographic Alignment Industry*, dissertation, Harvard University, 1988), who found the new-architecture aligners produced by the established manufacturers to be inferior in performance to those produced by entrant firms. One possible reason for these different results is that the successful entrants in the photolithographic aligner industry studied by Henderson brought to the new product a well-developed body of technological knowledge and experience developed and refined in other markets. In the case studied here, none of the entrants brought such well-

developed knowledge with them. Most, in fact, were *de novo* start-ups composed of managers and engineers who had defected from established drive manufacturing firms.

13. This finding is similar to the phenomenon observed by Joseph L. Bower, who saw that explicit customer demands have tremendous power as a source of impetus in the resource allocation process: "When the discrepancy (the problem to be solved by a proposed investment) was defined in terms of cost and quality, the projects languished. In all four cases, the definition process moved toward completion when capacity to meet sales was perceived to be inadequate. . . . In short, pressure from the market reduces both the probability and the cost of being wrong." Although Bower specifically refers to manufacturing capacity, the same fundamental phenomenon—the power of the known needs of known customers in marshaling and directing the investments of a firm—affects response to disruptive technology. See Joseph L. Bower, *Managing the Resource Allocation Process* (Homewood, IL: Richard D. Irwin, 1970) 254.
14. In booking \$113 million in revenues, Conner Peripherals set a record for booking more revenues in its first year of operation than any manufacturing company in United States history.
15. This finding is consistent with what Robert Burgelman has observed. He noted that one of the greatest difficulties encountered by corporate entrepreneurs has been finding the right "beta test sites" where products could be interactively developed and refined with customers. Generally, a new venture's entrée to the customer was provided by the salesperson representing the firm's established product lines. This helped the firm develop new products for established markets but not to identify new applications for new technology. See Robert A. Burgelman and Leonard Sayles, *Inside Corporate Innovation* (New York: The Free Press, 1986) 76–80.
16. I believe this insight—that attacking firms have an advantage in disruptive innovations but not in sustaining ones—clarifies, but is not in conflict with, Foster's assertions about the attacker's advantage. The historical examples Foster uses to substantiate his theory generally seem to have been disruptive innovations. See Richard J. Foster, *Innovation: The Attacker's Advantage* (New York: Summit Books, 1986).

CHAPTER TWO

Value Networks and the Impetus to Innovate



From the earliest studies of the problems of innovation, scholars, consultants, and managers have tried to explain why leading firms frequently stumble when confronting technology change.

Most explanations either zero in on managerial, organizational, and cultural responses to technological change or focus on the ability of established firms to deal with radically new technology; doing the latter requires a very different set of skills from those that an established firm historically has developed. Both approaches, useful in explaining why some companies stumble in the face of technological change, are summarized below. The primary purpose of this chapter, however, is to propose a third theory of why good companies can fail, based upon the concept of a *value network*. The value network concept seems to have much greater power than the other two theories in explaining what we observed in the disk drive industry.

ORGANIZATIONAL AND MANAGERIAL EXPLANATIONS OF FAILURE

One explanation for why good companies fail points to organizational impediments as the source of the problem. While many analyses of this type stop with such simple rationales as bureaucracy, complacency, or