CHAPTER 1

GETTING TO THE FRONT OF THE PACK

Each year, my wife, Miriam, our kids, Hannah, Eve, and Jesse, and I watch the Boston Marathon, which passes near our home. After the cacophony of the police escort and the press teams roaring past, there is a surreal calm as the first one or two runners fly by. Nearly two hours into the race, with just three miles to go, their form is flawless, their breathing easy, their faces calm. Then the clamor resumes.

A few dozen yards behind the leaders is a tight knot of athletes, all world-class but not looking as good. Their rhythm is a little off; their expressions are slightly pained. They are jostling and elbowing each other, but for all the effort, their only hope is to be runner-up, chasing the front-running, pacesetting leaders who are pursued but never caught.

The Boston Marathon only happens once a year, but every day we can see the same kind of ferocious competition among companies fighting for a consolation prize while one or two firms cruise to a victory which appears to be easy. In automobile manufacturing, commercial aviation, metal processing,

integrated-circuit fabrication, financial services, and health care, just to name a few, we can find "fair" contests in which opponents go head to head in the same product categories, woo the same customers, source from the same suppliers, hire from the same labor pools, struggle with the same dangerous conditions, and obey the same regulations. The playing fields are so level and there is so little differentiation among the rivals that one should expect cutthroat, tooth-and-nail, dog-eat-dog competition, fleeting profitability, and unsustainable leadership. And for many companies, that's how it is. Yet a few leaders are way out ahead, chased but never caught, generating a greater range and a higher quality of products and services, responding more quickly to the changing market, with fewer people, fewer resources, and fewer mishaps and accidents. While everyone else struggles to keep up, these *high-velocity organizations* race from success to success with growing market share, profitability, and reputation. In the marathon, everyone starts together and everyone crosses the half-way and three-quarters marks. The critical difference, of course, is that the leaders hit each milestone first and, by the time their challengers get there, they are well on their way to the next one. So it is among organizations, as represented in Figure 1-1. Everyone advances over time, improving performance along various metrics such as quality, efficiency, product or service variety, workplace safety, and time to market. The problem for the pack is that the market leader achieves a certain level before everyone else and, while others close in on where the high-velocity leader was, it has darted away, still to be chased but not captured.

Figure 1-1 High performance through superior improvement, innovation, and invention

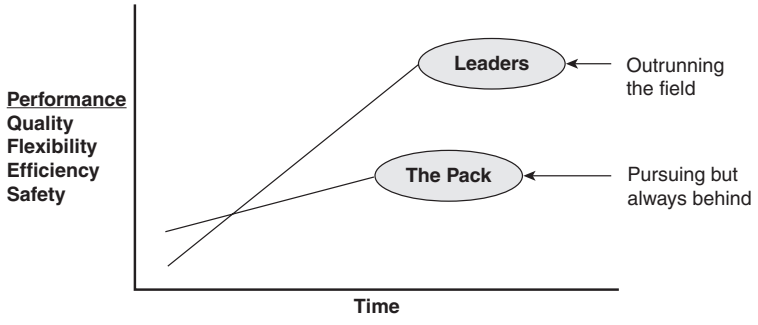

High-Velocity Organizations Abound

Let me offer a few examples, beginning with the automobile industry. Every major manufacturer makes cars, trucks, SUVs, and minivans. Those vehicles come in economy, regular, and luxury versions and in small, medium, and large sizes. The manufacturers contend for customers in every major market; their dealerships are often within walking distance of each other. They have design and production facilities in every region, hire in all those places in overlapping job markets, and are subject to the same regional rules and regulations. They often buy from the same suppliers. I worked in a plant with people making parts for Toyota while many of the same people, using the same equipment, were also making parts for direct competitors.

In this highly competitive environment, while General Motors (GM) and Ford struggle from one year to the next and Daimler has shed Chrysler after destroying tens of billions of

dollars in shareholder value in an ill-fated merger, Toyota roars from success to success. It raced past General Motors as the world's production leader, ran by Ford to become the second-largest seller of automobiles in North America, and passed Chrysler as the third-largest automaker in North America. While Ford shed its luxury brands, Toyota's Lexus, a relatively recent entrant, pushed ahead to become the best-selling luxury brand in the United States. The Scion, an even newer introduction, is accomplishing what has proved to be difficult for other automakers: attracting young buyers to an established maker. Despite long-standing claims by competitors that high-mileage, high-performance, low-emissions cars are a technological and financial impossibility, Toyota launched the Prius, built market share, and bested its counterparts in establishing a standard for hybrid-drive technology, which now is found across its product line. While most auto companies were shutting plants and laying off employees, Toyota expanded, creating more opportunity to widen the gap further.

All this has led to staggering profitability. Toyota crossed the $10 billion threshold in 2003. In the fiscal year ending March 2007, its net income was $13 billion, compared with losses of $2 billion and $12.6 billion at GM and Ford, respectively. Toyota's market capitalization of $187 billion was greater than that of GM, Ford, and DaimlerChrysler combined. And all this occurred despite the fact that Toyota entered the U.S. market with few products, little brand-name recognition (and even less that was positive), and no manufacturing facilities decades after its competitors were well established.

Toyota is not alone in setting itself apart in a tightly competitive market. In commercial aviation, every major airline

buys equipment from the same vendors: Boeing and Airbus for large planes; Saab, Embraer, and Bombardier for regional jets; and General Electric, Rolls-Royce, and Pratt & Whitney for engines. Jet fuel is a commodity. The airlines use the same labor pool for pilots, flight attendants, gate agents, baggage handlers, and mechanics, and they compete for exactly the same customers flying between the same cities. This makes it hard for most carriers to differentiate themselves, with predictable results. Year in and year out, American, United, USAir, and the others face financial difficulties, demanding concessions from their workforces and expecting customers to put up with less comfort, worse service, and reduced reliability.

This is not so, however, with Southwest. Achieving a combination of low cost and high customer satisfaction, this airline has generated an annual profit for more than 30 years in a row, despite the spikes in fuel prices, declines in travel after 9/11, overcapacity in the industry, and price cutting by incumbents trying to fend off entrants. Whereas the industry as a whole has had a 50 percent loss in stock market value in the last decade, Southwest's valuation has doubled. Even since 9/11, Southwest has fared better than its competitors, with only a 20 percent drop in value versus 70 percent for the entire segment.

Consider another way to measure Southwest's disproportionate success in its market: In fiscal year 2006, the combined revenue for American, Continental, Delta, JetBlue, United, US Airways, and Southwest was $95.2 billion, of which Southwest accounted for 10 percent. In November 2007, the combined market capitalization of those airlines was $33 billion, of which Southwest accounted for 33 percent.

How has this been possible? According to my colleague Jody Hoffer-Gittell and others, some of the intuitively obvious answers are wrong. Southwest is as unionized as the other airlines, it has competition on all its routes, and it doesn't have the advantages of monopolistic pricing that the hub-and-spoke system gives the major carriers over some routes. So it is not succeeding thanks to some structural advantage. Rather, Southwest does the basic work of running an airline better than other airlines do—turning its planes around at the gate in less time with less effort and greater predictability and performing scheduled maintenance with greater reliability. Its crews and equipment therefore spend more time aloft with paying customers rather than sitting on the ground unprofitably and unproductively.

Manufacturing integrated circuits—microprocessors, memory chips, application-specific integrated circuits—can be brutally competitive. All "fabs," as the manufacturing facilities in this industry are called, buy equipment from the same vendors, make products that compete on the same dimensions of "device density" and speed, and sell them to the same electronics companies. Yet in this business too, some companies outrace their rivals. According to the Competitive Semiconductor Manufacturing Program at the University of California at Berkeley, there are significant disparities among competitors in terms of the performance levels they achieve for quality (e.g., defects and yields), speed (e.g., throughput and cycle time), and efficiency (e.g., labor productivity) and also, more notably, the speed with which those levels are achieved (e.g., process-development time and ramp-up time). Christensen, Verlinden, King, and Yang, in their article "The New Eco-

nomics of Semiconductor Manufacturing," give an example of how this comes about. They detail how one anonymous manufacturer, through an intense focus on operational excellence, cut the manufacturing time for a wafer by two-thirds and the cost per wafer by 12 percent. Effective capacity went up 10 percent and the number of products the plant could sustain increased by half. This plant became faster at meeting a broader range and volume of demand at a lower cost and with no extra capital investment.

Alcoa is in the business of mining, refining, smelting, forging, casting, rolling, and extrusion—all of which are inherently dangerous processes. Yet, during the late 1980s and early 1990s, a period of great business success for Alcoa, it established itself as the safest large manufacturing employer in the United States. According to recent Occupational Safety and Health Administration (OSHA) data, Alcoa's workplace injury rate is one-quarter the average for all manufacturers by one measure and one-twentieth by another. This wasn't accomplished by any competitive maneuvering. Something else enabled Alcoa to just say no to work-related accidents. How this has been accomplished is explored in detail in Chapter 4.

Not all high-velocity organizations are running for profit. Some measure performance in other ways. For example, nearly all leading hospitals have access to cutting-edge science, the latest technology, and intelligent, well-trained, hardworking, well-meaning employees. Yet there are large variations in safety. On the whole, hospitals are dangerous places for patients. The Institute of Medicine estimated that up to 98,000 of the 33 million Americans who are hospitalized each year die because something went wrong in the management of their

care. Other studies estimate that an equal number die as a result of an infection acquired while hospitalized and that an even greater number are nonfatally injured or infected in the course of receiving care. This puts the risk of suffering harm while being hospitalized as high as one in a few hundred and the risk of being killed as high as one in a few thousand. Yet a few hospitals have cut the risk that patients will be harmed by medical error and infections by 90 percent and more, putting themselves in a position to provide far better care to more people at less cost and with less effort than is typical elsewhere. These hospitals, like Alcoa, have that special "something else."

Being a crew member on board a nuclear-powered submarine might seem a risky proposition, as it might mean sharing space with nuclear-tipped warheads, with your ship subject to crushing pressures, while playing cat and mouse with adversaries' warships, all while operating blind and sometimes deaf. And we all have our impressions of nuclear energy, given the events at Chernobyl and Three Mile Island.

However, nuclear-powered warships in the United States Navy have collectively accumulated over 134 million miles and over 5,700 reactor-years of nuclear reactor operation since the first nuclear-powered submarine, the *USS Nautilus*, was launched in September 1954. In all that time, with all that use, there has not been a single reactor-related casuality or fatality. In contrast, the Russian nuclear navy has been far more accident-prone. NASA, also charged with manned missions in a hostile environment, has had a tarnished record. We'll take a closer look in Chapter 3 at why NASA has been problem-plagued and, in Chapter 5, will contrast this with the Navy's approach.

High-Velocity Competitors

What is the special "something else" that separates high-velocity organizations from their rivals? There is a rich research history of attempts by practitioners and academics to answer that question. Let's look at that history to better understand what *The High-Velocity Edge* contributes.

By the 1980s, the post–World War II political and military rivalry between the United States and its allies and the Soviet Union and its allies, which had demanded so much attention for decades, was finally quieting down. However, all was not smooth sailing. An increasingly wide array of formerly stalwart American industries and corporations faced a severe competitive threat. Foreign companies, many of them Japanese, were delivering higher-quality products at lower costs than seemed possible. The implications for America's economic well-being were staggering.

Initially, this phenomenon was explained in terms of economic conflict, perhaps because the Cold War mind-set still prevailed. Books such as Chalmers Johnson's *MITI and the Japanese Miracle* (1982) and Clyde Prestowitz's *Trading Places: How We Allowed Japan to Take the Lead* (1988) attributed Japan's success to a clever trade strategy masterminded by governmental ministries and coordinated with corporate networks (*keiretsu*) that outpaced the disjointed efforts of American companies, federal agencies, and Congress. According to this view, Japan rigged the game with advantageous financing structures, freedom from the pressures of what were characterized as shortsighted American financial markets, and a compliant population willing to delay gratification and suppress individual

interests to achieve corporate and national interests. It was a samurai culture versus a cowboy one, and with competitiveness defined as a contest among nations, the proper response to such "cheating" was thought to be national in scope: voluntary export restraints, domestic-content requirements, and industry-wide research consortia.

Inspired by that sort of explanation, I wrote my undergraduate thesis at Princeton on the macroeconomic determinants of exchange rates with the idea that understanding why the dollar was strong and the yen was weak might offer insights into ways to reverse the flow of goods and services. After college, my work in investment banking in the mid-1980s reinforced the notion of national economic competition. My colleagues and I were attuned to "what the Japanese would do" every time a new auction of government bonds took place. Later, working in Washington, D.C., for a congressional agency, I had a close view of the debates about restoring American competitiveness, which often focused on legislative and executive branch responses to such perceived infringements as subsidization and trade dumping.

Arriving at MIT as a graduate student in the late 1980s was fortuitous for me. The prevailing view of Japanese commercial ascendancy was shifting from a Cold War-style national competition to the management practices of individual market-leading firms. Books such as *Kaisha*, *Made in America*, *Dynamic Manufacturing*, and *The Machine That Changed the World*, along with a slew of articles, detailed the differences in business practices—particularly in design and production—between the new Japanese winners and the American firms they were displacing. This shift in emphasis proved to be extraordinarily productive.

It was observed that, at winning Japanese factories, products advanced to completion along simpler process flows than they did in American factories. Production was "pulled," triggered by actual customer need, rather than "pushed" in accordance with preconceived schedules. Work sites were more orderly and were organized according to the specific task that had to be accomplished at each location. Relationships with employees and suppliers tended to be collaborative, a far cry from the antagonistic industrial relations in America.

Also observed was the relentless *kaizen* (improvement), a process of engaging those closest to the direct work of the organization in the continual improvement of that work. So it was not just the velocity of material through the factory that mattered; it was the velocity of improvement and problem solving—the speed with which these factories discovered problems and solved them.

Researchers such as David Garvin documented differences in productivity among similar plants and found discrepancies of tenfold and even a hundredfold in quality. John Krafcik documented extraordinary differences in productivity between mass producers and lean producers in the auto industry. Michael Cusumano provided a historical account of Toyota's rise to ascendancy. James Womack, Dan Roos, and Dan Jones illustrated some of the major differences in shop-floor management, product design, and supplier relations between the auto industry's best and the rest in their landmark book, *The Machine That Changed the World*. John Paul MacDuffie revealed some of the details of the powerful problem-solving mechanisms these manufacturers employed.

Bob Hayes and Steve Wheelwright, with coauthor Kim Clark, put aside their focus on strategic decisions as the means toward

Restoring Our Competitive Edge and later wrote glowingly about the advantages of creating "the learning organization" in order to achieve world-beating *Dynamic Manufacturing*. Collectively, these and other authors conveyed the palpable sense of urgency found throughout the market-leading organizations to identify market needs, meet those needs, and get ever better at doing so.

This new perspective was exciting. It meant that managers mattered. Even if a firm's external environment was hostile, its internal environment could be shaped to positive effect. Managers did not need government to rescue them, nor did they have to skulk around the marketplace looking for arenas bereft of competitors. They could do what the Japanese were doing and take them on in a fair fight.

Inspired by these discoveries, many people, my classmates in the MIT-Japan Program and I included, threw ourselves into understanding Japanese management so that we could do our part in helping the United States recover from its competitive malaise. Many of us joined Japanese companies for an insider's view. For me, this meant dipping my toes in the water of Japanese business at a commercial bank in the summer of 1990 through the support of the Japan Society of New York and the International House of Japan (Tokyo) and then spending more than a year as part of an international manufacturing consortium at the University of Tokyo with the support of the Japanese Ministry of Education. I worked with Japanese, Germans, French, and Canadians from construction firms, industrial equipment manufacturers, and electronics companies, all of whom were trying to understand what their firms had to do in the face of accelerated technological innovation and heightened cross-border trade and competition.

When I returned to the United States in the mid-1990s, I noticed something strange. The groundbreaking research cited above, which had shown the enormous disparities between the best in an industry and the rest, was now nearly a decade old. In that interval, Toyota, the company that epitomized the Japanese approach (which by then had come to be called "lean manufacturing"), had been studied relentlessly. Hundreds of thousands of visitors had toured its NUMMI joint venture with General Motors in Fremont, California, and its greenfield site in Georgetown, Kentucky. Countless pages had been written about Toyota specifically and lean manufacturing more generally. Hundreds of manufacturing companies had benchmarked the company and each of the American Big Three had created its own version of the Toyota Production System (TPS): the Ford Production System, the Chrysler Operating System, and the GM Global Manufacturing System. All over, people were mastering the intricacies of pull systems, work standardization, and the like, yet no American Toyota had emerged.

Here was the problem: Although Toyota's competitors had indeed improved in both initial quality and manufacturing efficiency, Toyota had not been sitting still. High-velocity organizations don't. Not only had it also improved in quality and efficiency, it had expanded the range of the competition. It had localized production, increased its product offerings, introduced new technology, and created new brands. I'm reminded of football: Everyone was trying to improve the running game, and then a few teams invented the passing game. As the other teams tried to add passing to their playbooks, the leaders put the receivers in motion and added quarterback options and

calling plays at the line of scrimmage, always complicating the challenge by increasing the speed of the game and the range of plays that might occur.

When I entered Harvard Business School as a doctoral student, I set out to learn why it was so hard to overtake Toyota, and in the next four years I had extraordinary opportunities to do just that. The heart of my studies was learning by doing. For six months I was part of a Toyota team, working to develop a first-tier supplier in Kentucky (the one mentioned earlier that also supplied two of Toyota's competitors) and learning the Toyota Production System firsthand by solving production-related problems and working with others to do that. To appreciate the differences between what we were doing at the supplier and how more traditional manufacturers operated, I prepared by spending a week doing assembly-line work at one of Toyota's American competitors. We'll see more of that experience in Chapter 3. To appreciate the management of work systems across a broad range of products, processes, markets, and regions, I traveled to three dozen plants in North America and Japan to make observations, collect data, and interview people, from frontline workers to plant managers and corporate executives.

What I found was completely unexpected. I had already studied what had been written about Toyota, lean manufacturing, Six Sigma, and total quality management. I had a fairly good conceptual understanding of work standardization, pull versus push, the design of experiments, statistical process control, and the many other analytical and control tools that were being popularized. I thought I was looking for a still-missing tool or two. I couldn't have been more wrong.

The difference between Toyota and its competitors was neither more tools nor more diligent application of tools that had gained wide currency. That approach promised gains that were potentially significant but that would ultimately plateau. Michael Porter made that point in his 1996 *Harvard Business Review* article, "What Is Strategy?" If everyone benchmarks the leader by imitating how work is done at a particular time and place, no one can do any better than the leader and everyone will look and act the same, commoditizing their sector and guaranteeing that no one will enjoy an advantage.

Rather, what I was coming to appreciate was an approach to managing exceptionally complex work that mustered the hands *and* minds of hundreds of people so that improvement, innovation, and adaptation were unending. The factory was not only a place to produce physical products, it was also a place to learn *how* to produce those products and—most important of all—it was a place to *keep* learning how to produce those products. In fact, this is exactly what so much of the early research about Japanese management had revealed—that learning and discovery were intrinsic to success. But that idea had gotten lost as people focused on the particular tools and artifacts used in the workplace at the expense of understanding the principles of how those systems were managed.

The emphasis on learning and discovery went right to the heart of a fundamental managerial challenge. Complex products and services require complex design, production, and delivery operations. Organizations need to master the myriad functions that have to be brought to bear, but that alone will never be sufficient. They also need to master the countless permutations with which the various people, parts, and processes

can interact within such complex product and service operations. Such mastery is never complete—it can never be designed into the operation from the start.

For example, the Toyota plants that I visited were enormous, some with hundreds of millions of dollars in equipment, dozens if not hundreds of managers, and hundreds if not thousands of hourly workers. One would expect such massive operations to have an unavoidable inertia, but my key impressions were of movement and change, much of it urgent and adrenaline-charged. This was true both for work by an individual—such as installing a seat in a car, attaching a bumper, or connecting wiring—and for complex work carried out by large groups—such as launching a new model or building a new plant. No matter what the task, Toyota had figured out how to do the work in such a way that individuals and groups kept learning how to do that work better. Good luck benchmarking that. Any snapshot would reveal where Toyota was today but not where it was headed. Later, when I began to seek out and explore other high-velocity organizations in other fields, I was to find several that had independently arrived at the same idea, strengthening my conviction that the approach described in *The High-Velocity Edge* will help any organization engaged in complex operations to improve its performance.

Though many firms had embraced various tools associated with lean manufacturing and total quality management and had gained stability and control of work sites that had been chaotic and unreliable, they still never caught up. And now I could see why. These firms had picked up the visible tools of high-velocity organizations—the value-stream maps, pull

systems, production cells, statistical process control charts, and design of experiments—but they had not understood what these tools were for: managing complex work for continual improvement of that work (and therefore of the products and services that result from that work). As Kent Bowen and I pointed out in our 1999 *Harvard Business Review* article, "Decoding the DNA of the Toyota Production System," copying the tools alone did not generate the paradoxical combination of stability and flexibility that was increasingly associated with Toyota. It was Toyota's way of designing and improving processes that generated both short-term stability and longer-term agility and responsiveness.

As my research at Toyota progressed, a marvelous opportunity arose to test my findings. Alcoa had been pursuing the audacious goal of creating a perfectly safe work environment, despite the hazards that seemed inherent in its production processes. It was coming pretty close. The key for Alcoa, as we shall see in Chapter 4, was to realize that perfect safety could not be designed into its work from the start. No brain trust could ever figure out in advance all the little things that could go wrong. Instead, the trick was to do work, take immediate notice of any risks or potential risks in the work, and make changes so that the same risks did not reappear. And finding one risk wasn't an isolated experience. Pulling on the thread revealed many other process shortcomings that had not been known. In the area of safety, Alcoa had begun developing a management system much like Toyota's, in which the creation of products and the operation of processes were coupled tightly with creating better methods for being successful. Although the perfect safety system could not be designed, it

could be discovered bit by bit if enough velocity were generated and enough energy were sustained.

But could this Toyota-like approach be applied to Alcoa's business as a whole, a business very unlike Toyota's? In short, did my Toyota findings apply only to Toyota and to similar industries, or were they much more broadly applicable? In 1997, I worked with a group at Alcoa to develop and deploy the Alcoa Business System, based on the Toyota Production System. Some of the results were fantastic, as we will see in Chapter 4.

But the circle was to widen again. In early 2000, there was a knock on my door at the Harvard Business School, where I was now on the faculty. In walked a doctor named John Kenagy. "I'm a vascular surgeon," he explained, "and my colleagues and I have tried everything we can to raise the quality and efficiency of our practices and of the hospitals in which we work. Nothing has helped. I've heard about this Toyota research you've been doing. Could a similar approach work in health care?"

We didn't know. Here, indeed, was another kind of very complex service being provided by a very complex organization and, as I was vividly to learn, working in a hospital can be a stressful experience with little failures happening all the time, some of which might prove dangerous or fatal to patients in unexpected ways. Could the often-frustrating work of nurses, aides, doctors, administrators, and staff be managed in a way that was dynamic, adaptive, self-improving, and self-innovating? We gave it a try, first at Deaconess Glover Hospital in Needham, Massachusetts, and later at a number of

hospitals through the auspices of the Pittsburgh Regional Healthcare Initiative. The results, some to be discussed in Chapter 11, were stupendous.

What do all these examples mean for you, the reader? I and other researchers have found—and in a few cases I myself helped create—high-velocity organizations engaged in a wide variety of missions. As different as these organizations are in many respects, they have one thing in common: They are adept at designing, developing, and operating exceptionally complex systems to achieve exemplary and constantly improving performance in the design, production, and delivery of complex goods or services. This is the "something else" that is needed when monopolistic advantage or a lower level of performance are not viable options. This is how the market leaders get ahead and stay ahead.

At this point, we have looked at the class of front-runners who are clearly doing something different than their peers and competitors, something that helps them take the lead and then keep increasing their lead. We have also asserted that it is not enough to imitate the distinctive techniques of these front-running leaders, to mistake the means for the ends. It is necessary to understand the goal of those techniques and to dedicate the organization's efforts to that goal—the management of complex operations for high performance.

But having given examples of high performance and having used a historical survey to clarify the real goal, I would like to say some more about the means.

Structure and Dynamics of High-Velocity Organizations

At a high level, we can distinguish two characteristics that distinguish high-velocity organizations from those struggling behind them.

1. Structure: Managing the Functions as Parts of the Process

There is a structural difference between the high-velocity organizations and those chasing them that creates potential for speed. While high-velocity organizations put great effort into developing the technical competency of various functions, they are equally and always concerned with the way the work of individuals, teams, and technologies will contribute to (or impede) the process of which they are part. The process orientation of high-velocity organizations is in contrast to the "siloization" of so many other organizations in which the departments may talk of integration but tend to operate more like sovereign states. In high-velocity organizations, functional integration is not just pretty talk, it is the nuts-and-bolts of management at all levels every day.

2. Dynamics: Continually Improving the Pieces and the Process

There is a dynamic difference between the high-velocity leaders and those chasing them that generates speed. High-velocity organizations are constantly experimenting and learn-

ing more about all the work they do; this is how they cope successfully with the complexity which they all face in one form or another. These organizations do not encourage or admire workarounds, firefighting, and heroic measures. They want to understand and solve problems, not put up with them.

It would be impossible to exaggerate how valuable this is. How much time and effort is saved by getting rid of a problem once and for all? How much confidence is gained when people see that they don't have to keep putting up with one problem after another and that management doesn't want them to? How many *more* problems will be solved because people know they can? Then there is the paradoxical benefit that solving one problem often reveals another that had been masked by the first one. Another problem, yes, but now the organization sees it as yet another problem *that's going to be gotten rid of.*

Low-performing, low-velocity organizations are strikingly different. First, they tend to be *functionally oriented* and do not manage the relationships among all the elements adequately, as was mentioned above. Second, even if they think in terms of processes, they are not dynamic. Instead of constantly doing work, watching for problems in their approach, and modifying the way they work, they lock into an approach that seems good at the time and—even when it proves inadequate—stick with it and muddle through.

To sum up, high-velocity organizations differ from low-velocity organizations both structurally and dynamically. Structurally, they insist that each piece of work be done with an eye to the larger process of which it is a part. Dynamically, they insist that each piece of work be done in such a way as to

bring problems to the attention of those who can best analyze and solve them. Low-velocity organizations, in contrast, are characterized by "siloization"—"You do your job and I'll do mine"—rather than integration and by endless workarounds and firefighting—"This'll do for now" or "Don't worry, this happens all the time"—rather than continual improvement, innovation, and invention.

The Four Capabilities of High-Velocity Organizations

The ability of high-velocity organizations to be so functionally integrated and continually self-improving, innovative, and inventive is rooted in four complementary capabilities. I will explain each of them briefly here. They will turn up again and again in Chapters 3 through 5 and they will be explored in detail in Chapters 6 through 9. Note that Capability 1 is the key to functional integration for high performance, while Capabilities 2 through 4 are the keys to managing an organization for continual self-improvement.

Capability 1: Specifying Design to Capture Existing Knowledge and Building In Tests to Reveal Problems

High-velocity organizations don't like anyone to start work, whatever its size or complexity, until the organization has (1) specified the most effective approach that is currently known for achieving success at that task and (2) built into that

approach the capacity to detect failure when and where it occurs.

Whether the work is to be done by an individual or a group, with or without equipment, high-velocity organizations are uncomfortable with ambiguity. They specify in advance what (a) outcomes are expected; (b) who is responsible for what work in what order; (c) how products, services, and information will flow from the person performing one step to the person performing the next step; and (d) what methods will be used to accomplish each piece of work.

However, it is not that they want or need guarantees. This kind of specification is not a case of perverse Taylorism or micromanagement, with smart people telling less-intelligent people what to do. It is, in fact, an investment. Before the work starts, the high-velocity organization invests everything it knows so far into these specifications to maximize the likelihood that people will succeed.

But this is the sort of investment that has a positive payout regardless of the immediate outcome. Specifying with clarity and care what actions are expected to lead to what outcomes makes it far easier to recognize when something unexpected has happened. This highlights gaps in the organization's collective knowledge about how to succeed. With pockets of ignorance identified, the high-velocity, front-running organizations know where they need to invest to get better. To increase their ability to discover what they don't know, they even go out of their way to build tests into their operations in order to detect abnormalities when and where they occur. In contrast, those laboring in the pack are less committed to upfront specification, already handicapping themselves from the

start, since they are not using the best possible approach. And then they suppress their ability to see when what they are doing is not good enough. Like an athlete who uses antiquated equipment and doesn't keep on eye on the competition, they find themselves falling farther and farther behind.

Capability 2: Swarming and Solving Problems to Build New Knowledge

High-velocity organizations are adept at detecting problems in their systems at the time and place of their occurrence. They are equally adept at (1) containing those problems before they have a chance to spread and (2) diagnosing and treating their causes so the problems cannot reoccur. In doing so, they build ever-deeper knowledge about how to manage the systems for doing their work, converting inevitable up-front ignorance into knowledge.

It all happens like this: In high-velocity organizations, problems are swarmed at the time and place where they occur and by the people who are affected. A benefit to swarming a problem immediately is that it can be contained before it can affect someone else's work. And the longer the problem remains unresolved, the more difficult and more expensive it will be to solve. In Chapter 3, we'll see examples of what happens when problems are left untreated.

Swarming a problem is not only beneficial in terms of what is prevented—an infectious spread of the problem's impact. It is beneficial in terms of what is allowed—the gathering of essential, contextual information that would otherwise be lost to fading memory and changing circumstances. Many prob-

lems occur because of some unexpected, idiosyncratic interaction of people, processes, products, places, and circumstances. As time passes, it becomes impossible to reconstruct exactly what was going on when the problem arose.

Once swarmed and investigated, problems are solved, but not in any ad hoc, willy-nilly fashion. High-velocity organizations insist that "the scientific method" be used in a disciplined fashion. This is not an esoteric, ivory tower exercise; it reflects the conviction that when something is changed, those making the alteration should have a clear idea of what actions are expected to lead to what outcomes and should then be able to see when they are right and wrong. Fixing the problem isn't good enough; they want to fix it while gaining a deeper knowledge of how their own processes work.

Before moving on to the third and fourth capabilities, let me point out that the first two alone are game-changing. Many people set out to do work and are either successful or not. If not, the effort was wasted. High-velocity organizations convert win-lose situations into win-win situations. If they succeed, they win. If they do not, they learn how to succeed next time, and that is also a win.

Capability 3: Sharing New Knowledge throughout the Organization

High-velocity organizations multiply the power of their new knowledge by making it available, not only to those who discovered it, but also throughout the organization. They do this by sharing not only the solutions that are discovered, but the processes by which they were discovered—what was learned

and how it was learned. While their competitors allow problems to persist and propagate into the larger system because the solutions, if they are found at all, remain contained where they were found, the high-velocity leaders contain their problems and propagate their discoveries. This means that when people begin to do their work, they do so with the cumulative experience of everyone in the organization who has ever done the same work. We'll see several examples of that multiplier effect.

Capability 4: Leading by Developing Capabilities 1, 2, and 3

Managers in high-velocity organizations make sure that a regular part of work is both the delivery of products and services and also the continual improvement of the processes by which those products and services are delivered. They teach people how to make continual improvement part of their jobs and provide them with enough time and resources to do so. Thus, the organization's ability to be both reliable and highly adaptive becomes self-reinforcing. This is a fundamental difference from their also-ran competitors. High-velocity managers are not in place to command, control, berate, intimidate, or evaluate through a contrived set of metrics, but to ensure that their organizations become ever more self-diagnosing and self-improving, skilled at detecting problems, solving them, and multiplying the effect by making the solutions available throughout the organization.

Certainly, the idea that success comes to those who learn the most quickly and effectively has antecedents and, before we move on, let's recognize some of those. After all, the point

of this book is not to refute that previous research, but to show that many of these ideas are actually part of a holistic approach to managing complex systems for great outcomes. For example, Nelson and Winter emphasize, in *An Evolutionary Theory of Economic Change*, that managers don't necessarily plan their organizations' way to greatness, but that successful organizations develop routines, test them in practice, recognize which don't work, and reinforce those that do. Eric von Hippel and his coauthors have demonstrated the importance of learning in context. Because there are so many circumstantial factors that cannot be codified, learning must occur when and where problems are experienced. My late colleague Jai Jaikumar had "information perishability" as one of his axioms of information. Information is not only contextual, it spoils; that is why it is so important to swarm problems. More than a few writers have emphasized that self-reflective experience is critical to improvement. This point is highlighted in Chapter 4 in the Alcoa example and later in the chapters that focus on Toyota.

Chapter Overview

The High-Velocity Edge is intended to help readers understand how market leaders outdistance the competition and how great companies can catch up and win. It does so in the following fashion:

In Chapter 1, I have introduced a category of "high-velocity organizations" whose ability to consistently outperform their competitors cannot be explained well by manipulation of their external environment—competitors, suppliers, regulators, investors, and so on. It is explained largely by their mastery of

their internal environments—the complex operations needed to produce or provide complex products or services. This mastery boils down to the four capabilities just described, all of which contribute to these organizations' ability to discover more quickly and to bring discoveries to bear in accomplishing the organization's mission.

Chapter 2 explores in more detail the basic challenge of complex operations which all high-velocity organizations face. The main point is that the very scientific discoveries that inspire or improve the products and services on which we depend also increase the difficulty of managing their design and delivery. We'll look more closely at how systems evolve from simple and linear to complex, highly intertwined, and strongly interconnected, and what challenges that presents. Supporting the premise that the themes of *The High-Velocity Edge* are independent of particular sectors, one example is from the design and production of a manufactured product, and the other is from medical care.

Chapter 3 is the "doom and gloom" portion of the book, in which we look at approaches to managing complex work that bring all kinds of frustration, waste, and failure, ranging from the time nurses spend looking for rubber gloves to the sudden demise of two space shuttle crews to the slow-motion failure of once-grand automotive corporations. While the contexts are different, the failure modes are nearly identical.

Things look up from there. Chapter 4 provides a detailed example (the first of several) of how exceptionally complex work can be managed for outstanding results. We'll see how Alcoa converted itself into the safest manufacturing employer in the country by shifting from an approach more typical of

the organizations in Chapter 3 to a dynamic discovery approach based on seeing problems, solving problems, and sharing quickly and broadly what was learned—all this supported by senior leadership.

Chapter 5 shows how the same commitment to managing systems with a bias toward discovery led to great success for several other organizations far afield from Alcoa and from each other. These are the U.S. Navy Nuclear Power Propulsion Program, Pratt & Whitney's jet engine design group, and Avenue A, an Internet advertising agency. As pointed out earlier, the variety of examples is evidence that we are talking about general principles, not the particulars of any one industry or setting.

Chapters 1 through 5 give an overview of the main thesis of *The High-Velocity Edge*, that some organizations achieve exceptionally high velocity in self-correction, self-improvement, and internally generated innovation and invention and use this velocity to set themselves apart in situations that should otherwise be intensely competitive or constraining. In Chapters 6 through 10, we'll look in depth at how one company, Toyota, puts the principles outlined above into action.

Chapter 6, after setting up Toyota as an example of a high-velocity organization, focuses on Capability 1—the design and operation of self-diagnostic systems. A simple, robust framework for describing processes will be introduced. Then we'll walk through several examples—from simple to complicated and from tangible to less so—showing how specification is used to help work start off strongly and how tests built into systems help catch problems before they metastasize.

Chapter 7 focuses on Capability 2—swarming problems to contain them and solve them. We'll see how several Toyota

teams learned how to solve problems and fix work processes so that the processes improved and, at the same time, the individual workers became more skillful and productive. We'll also see the same problem-solving discipline practiced at senior levels.

Chapter 8 is about Capability 3—how local discoveries are made useful throughout an organization. Common themes will emerge from an example of disseminating the most effective known methods of "master craftsmen," an example of capturing knowledge and using it over several product design cycles, and an example of collaborative problem solving and process improvement. The most compelling theme is that when the solution to a problem is discovered, the discovery process itself must be conveyed along with the solution.

In Chapter 9, we will turn our attention to the critical role of leaders in high-velocity organizations—their exercise of Capability 4. Like other leaders, they are responsible for setting objectives and allocating resources, but they are also the stewards of the three other capabilities by which organizational velocity is generated. They must deliver those capabilities to those for whom they are responsible.

Chapter 10 concludes our in-depth look at Toyota by showing how the four capabilities are brought to bear in crisis-recovery situations like the overnight loss of a critical supplier or the closure of an essential port of entry. Those people who hold the belief that the high-velocity approach applies only to repeatable processes and fosters only incremental improvements will see that it can produce results at a speed and on a scale that are astonishing to most.

With Chapter 11, we leave Toyota and turn to the important task of creating high-velocity organizations in the Amer-

ican health-care industry. Those in the health-care field will see that better care does not have to come at greater cost, nor do spending caps necessarily require denial of care. Other readers will see that the four capabilities can work wonders not only in capital-intensive, technology-driven sectors, but in knowledge-intensive, service-based, nonrepetitive situations.

Chapter 12 will tie some parting thoughts together as a conclusion.

Before Chapter 2 begins, I want to say again how privileged I have been to be exposed to the great organizations and people represented in this book and to the many others for whom there was not space. I've learned a great deal from them, enjoying the experience every step of the way. I hope that I allow you, the reader, to enjoy the journey and its discoveries as well.