

**“Spear has dazzled readers with his insights.”**

*—Harvard Business Review*

Winner of the  
**AMERICAN SOCIETY  
OF QUALITY'S  
CROSBY MEDAL**

# THE HIGH-VELOCITY EDGE

**HOW MARKET LEADERS LEVERAGE  
OPERATIONAL EXCELLENCE  
*TO BEAT THE COMPETITION***

**STEVEN J. SPEAR**

**FIVE-TIME SHINGO PRIZE AWARD WINNER**

**FOREWORD BY CLAYTON M. CHRISTENSEN**

**BESTSELLING AUTHOR OF *THE INNOVATOR'S DILEMMA***



# THE HIGH-VELOCITY EDGE

HOW MARKET LEADERS LEVERAGE  
OPERATIONAL EXCELLENCE  
*TO BEAT THE COMPETITION*

STEVEN J. SPEAR



New York Chicago San Francisco Lisbon London  
Madrid Mexico City Milan New Delhi San Juan  
Seoul Singapore Sydney Toronto

Copyright © 2009 by Steven J. Spear. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without prior written permission of the publisher and the author.

1 2 3 4 5 6 7 8 9 0 DOC/DOC 1 9 8 7 6 5 4 3 2 1 0

ISBN: 978-0-07-174141-5

MHID: 0-07-174141-0

This book was previously published with the title of *Chasing the Rabbit*.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional person should be sought.

—From a Declaration of Principles Jointly Adopted  
by a Committee of the American Bar Association  
and a Committee of Publishers and Associations

McGraw-Hill books are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. To contact a representative, please visit the Contact Us pages at [www.mhprofessional.com](http://www.mhprofessional.com).

This book is printed on acid-free paper.

# CONTENTS

PREFACE	vii
FOREWORD	xiii
ACKNOWLEDGMENTS	xxi
<b>CHAPTER 1</b>	
GETTING TO THE FRONT OF THE PACK	1
<b>CHAPTER 2</b>	
COMPLEXITY: THE GOOD NEWS AND THE BAD NEWS	33
<b>CHAPTER 3</b>	
HOW COMPLEX SYSTEMS FAIL	45
<b>CHAPTER 4</b>	
HOW COMPLEX SYSTEMS SUCCEED	87
<b>CHAPTER 5</b>	
HIGH VELOCITY UNDER THE SEA, IN THE AIR, AND ON THE WEB	109



## CONTENTS

<b>CHAPTER 6</b>	
CAPABILITY 1: SYSTEM DESIGN AND OPERATION	155
<b>CHAPTER 7</b>	
CAPABILITY 2: PROBLEM SOLVING AND IMPROVEMENT	193
<b>CHAPTER 8</b>	
CAPABILITY 3: KNOWLEDGE SHARING	225
<b>CHAPTER 9</b>	
CAPABILITY 4: DEVELOPING HIGH-VELOCITY SKILLS IN OTHERS	263
<b>CHAPTER 10</b>	
HIGH-VELOCITY CRISIS RECOVERY	295
<b>CHAPTER 11</b>	
CREATING HIGH-VELOCITY HEALTH-CARE ORGANIZATIONS	323
<b>CHAPTER 12</b>	
CONCLUSION	357
EPILOGUE	365
REFERENCES	367
NOTES	377
INDEX	395



## PREFACE

**T***he High-Velocity Edge* shows how to beat your rivals, even in the most competitive markets. True, business can be brutal. You exhaust yourself identifying market needs, hustling to develop offerings, and rushing to produce and deliver those items. No sooner is this done than someone else competes for your customers, conspires with your suppliers, and quickly bids down returns on your efforts.

Still, a few enjoy a better experience. They too have to identify needs, create solutions, and deliver to market. However, rivals chase them and never catch up. Instead, the enviable few keep hitting new milestones. They win, not from finding pre-dation-safe positions. There are none. Rather, they discover ways to be better at what they do and develop products and processes that are better than anyone else's, operating with such velocity that pursuers are frustrated.

*The High-Velocity Edge* reveals how this is done. It shows how leaders create and sustain unassailable rates of broad-based, internally generated improvement, innovation, and invention, and how you and your organization can do so as well. And if



you succeed in transforming your organization into one that is high velocity, you too can face the market with unmatched combinations of reliability and responsiveness.

That you must compete not just by position but by velocity as well is the result of both external and internal pressures. Externally, you are always challenged by rivals in the private sector and increasingly by resource-strapped “patrons” in the public sector. Unless your organization figures out how to do more of what is better with less, you lose relevancy.

There is also internal pressure for high-velocity discovery, learning, innovation, and invention. The work of organizations is incredibly complex, and nothing complex can be designed perfectly. There are simply too many “parts”—be they the actual components of complex technical systems or the contributions of people who are expert in myriad disciplines—connected in too many convoluted, interdependent ways for a small group of smart people to plan a system that will work adequately—let alone perfectly. Since you cannot plan perfection, then you must pursue perfection.

*The High-Velocity Edge* develops the principles, with examples of their application, that leaders employ to ensure that their organization performs better than anybody else’s. This book explains how to:

- Create an organization in which work is done in a way that both harnesses the best known approaches currently available and signals loudly when and where new knowledge is needed.
- Solve problems as they arise, not only to make the immediate symptom disappear but also to develop new understanding that prevents the problem from recurring.





- Magnificently multiply the impact of these local discoveries by making them systemically useful.
- Lead an organization in a manner that facilitates the constant development of great discoverers thereby enabling discovery to happen all of the time, everywhere.

The organizations in *The High-Velocity Edge* differ widely by market, core technologies, and professional disciplines. Yet they have much in common. For example, they share responsibility for managing complex systems, and they also share a specific approach for creating organizations capable of unmatched learning and adaptation.

- Pratt & Whitney compressed the time and cost of designing jet engines despite the hurdles of developing and harnessing expertise in materials, air flow, combustion engineering, heat transfer, a jumble of manufacturing specialties, mechanics, electronic sensing and controls, and so on.
- Hospitals profoundly improved the quality, cost, and capacity of their care. Were similar achievement replicated broadly, we wouldn't be facing a national healthcare crisis.
- Toyota, once a crummy car maker in the late 1950s, raced ahead on productivity, quality, and reliability before expanding its model mix, creating new brands, regionally diversifying, and leading with hybrid drive technology. (We'll also deal with Toyota's recent struggles.)
- The U.S. Navy had one of the most successful "new product introductions" ever. In only six years, it created nuclear propulsion, which required developing new science and whole industries. Since launching the *USS Nautilus* in



1954, the Navy has had people running dangerous, complex machines in dangerous, complex situations—with perfect safety. *The High-Velocity Edge* examines the leadership that made such success possible.

- aQuantive, an Internet advertising pioneer, innovated its way from being a good but money-losing idea to a multi-billion-dollar enterprise.

One parting thought: Although the first version of *The High-Velocity Edge* was written before the financial crisis of 2008–2009, it is directly relevant amidst the current economic turmoil in two ways. One is in helping understand why failure occurred and the other is in providing guidance for recovery.

As for the first, you'll find *The High-Velocity Edge* an optimistic book in that it offers straightforward principles that, if practiced with discipline, lead to great outcomes. The stories inside are not all happy ones, however. Chapter 3, in particular, shows how the mismanagement of complex organizations leads to failure. The featured organizations may differ by product, service, and sector—financial institutions, storied manufacturers, and the like. However, many, if not most, failed for exactly the reasons explored inside—not seeing problems before they metastasized uncontrollably, and not solving problems even if they were seen with sufficient lead-time. Toyota clearly tumbled into trouble when its ability to grow the capabilities on which it depended for improvement and innovation was overburdened by the demands of rapid expansion. If anything, this emphasizes the importance of the capabilities explored in *The High-Velocity Edge* and explains why constantly nurturing them is so critical.



As for recovery, the importance of learning and adaptation cannot be overstated. During good times, organizations may have answered well enough the following questions: (1) what does the market need?, (2) what products or services can I offer?, and (3) how do I produce and deliver them adequately? Now, however, whatever answers you had to those questions before the financial crisis are no longer adequate. New needs demand new offerings that demand new ways of bringing them to market.

We've already seen organizations unable to generate new solutions to new problems. They've fallen by the wayside. More will follow.

We'll also see organizations that might not have been leaders going into the recession accelerating now and emerging ahead. I hope the lessons in *The High-Velocity Edge* help your organization become one of those.

Thank you for reading my book. Please let me know what you think.

With best wishes,

Steven J. Spear

Brookline, MA

February 2010

Web site: <http://www.TheHighVelocityEdge.com>





## ACKNOWLEDGMENTS

With discovery as a theme of *The High-Velocity Edge*, it is appropriate that I thank those who invested so heavily in the discoveries reported in this book. First is Professor H. Kent Bowen of Harvard Business School and of MIT before that. Kent opened the door for me to pursue graduate studies at the doctoral level, and his involvement in what I was doing went far beyond what anyone could expect or deserve on the strength of his or her own merits alone.

Hajime Ohba of Toyota has also been my teacher since I began this research. Toyota is one of the most intensively studied organizations in this era, with thousands of academic researchers, journalists, and practitioners seeking the firm's time and access. Toyota did not need another person pestering it for access, but Ohba invited me in nevertheless. While I was there, he made sure I learned without hesitation, obstacle, or obstruction. Whenever we were together, and there were many times since 1995, he allowed me to learn from his experiences and the perspectives he had developed, or he coached me toward discoveries of my own. Typically, it was both at the same time.



## ACKNOWLEDGMENTS

When I was a student and later a faculty member at Harvard Business School, Professors Carliss Baldwin and Clayton Christensen were constant with their inspiration and support. Those familiar with their work will see Carliss's influence on my treatment of the structure of complex systems and Clay's influence when I write about the dynamics of innovative competitors. Less visible but equally vital was their generosity in helping me express my ideas as they were taking shape. It is no small honor that Clay agreed to write the foreword for this book, once again contributing to the betterment of my work.

Learning is inherently collaborative, especially learning about how groups collaborate for maximum effect. I am indebted to thousands of people who have let me into their worlds and the daily work they do. They include but certainly are not limited to Tosh Akioka, Jamie Bonini, Toshi Kitamura, Olivier LaReau, Lisa Nichols, Christine Parker, Bryant Sanders, Cindy Voss, and others at Toyota and its suppliers. Keith Turnbull at Alcoa was a tireless colearner and teacher. I owe thanks to them and his many colleagues, including Arnaldo Cruz, Pat Love, John Marushin, and Stan Vishnevsky.

We did not know if lessons from great companies such as Toyota and Alcoa would apply to something as far afield from automation and heavy industrial processes as health care. The answer wasn't obvious. Jim Reinertsen and John Dalton at the Beth Israel Deaconess Medical Center created opportunities for my colleagues and me to test ideas in practice, and former Alcoa CEO Paul O'Neill made further exploration possible with the Pittsburgh Regional Healthcare Initiative, where we worked with many people, including Vickie Pisowicz, Mark



Schmidhofer, and David Sharbaugh. Paul's relentless commitment to the well-being of other people—at Alcoa, in Pittsburgh hospitals, and more generally—is inspirational.

David Champion, my editor at *Harvard Business Review* for ten years, has helped me bring into focus what was of the greatest interest and impact. John Elder has shaped the content and delivery of my work by playing the role of thoughtful lay reader and editor par excellence for papers, cases, and especially for this book.

Thank you to former MIT Sloan School dean Lester Thurow, who made problems of competitiveness, globalization, and technological progress a priority for a generation of students, and who believed that we could be part of the solution. Thank you also to Professor Mike Cusumano for his deep insights about great companies; to Professor David Hardt, who taught me so much about complex systems; and to Dick Samuels for the MIT-Japan Program.

I would be remiss not to thank my parents, grandparents, and great-grandparents for the examples they provided and the standards they set. Although any family will endure trials across several generations, mine has always demonstrated the importance of appreciating and repaying life's triumphs through kindness and service. I hope their legacy shows in my own work, as it certainly does in my brother Jonathan's efforts to advance innovative educational programs.

Finally, I have to say thank you close to home. Writing a book is a luxury, a rare opportunity to explore a topic in its many nuances. The time to do so is a gift. For this, I owe an unpayable debt to my wife, Miriam, and our children, Hannah, Eve, and Jesse. They endured my time in the field to learn



## ACKNOWLEDGMENTS

from other people and my late nights and early mornings of writing at home.

But the debt is deeper than that, as my research is inexorably intertwined with personal milestones—meeting Miriam the first week of school, courting while in the midst of field work, and celebrating the birth of our first child (nearly) at graduation. Moreover, *The High-Velocity Edge* emphasizes that great things come from sweating the details. Miriam models that as an architect, mother, and wife, and it is an example taken up by our kids. I would be immeasurably worse off without her and them.

*Steven Spear*  
*Brookline, MA*





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, pace-setting 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 casualty 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. I 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 “silozation” 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 “silozation”—“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 up-front 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 an 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.







## CHAPTER 4

# HOW COMPLEX SYSTEMS SUCCEED

**W**e now leave the failures behind to look at the successes. In a broad variety of sectors there are organizations with a much more productive approach to managing the complex operations on which they depend. Unlike their counterparts, who manage functional specialties in isolation from each other, without a view of the pieces in relation to a larger whole, the leaders invest continually in the integration of specialties into a process. Unlike their counterparts who dismiss the regular chatter of imperfect processes (and products) as unavoidable noise, they continually advance their expertise. When their operations speak up—in the language of problems or unexpected outcomes—these organizations stop, listen, learn, improve, and innovate, propagating what is learned in one situation to have maximum impact throughout the organization.



## Alcoa: Safety In Unsafe Situations

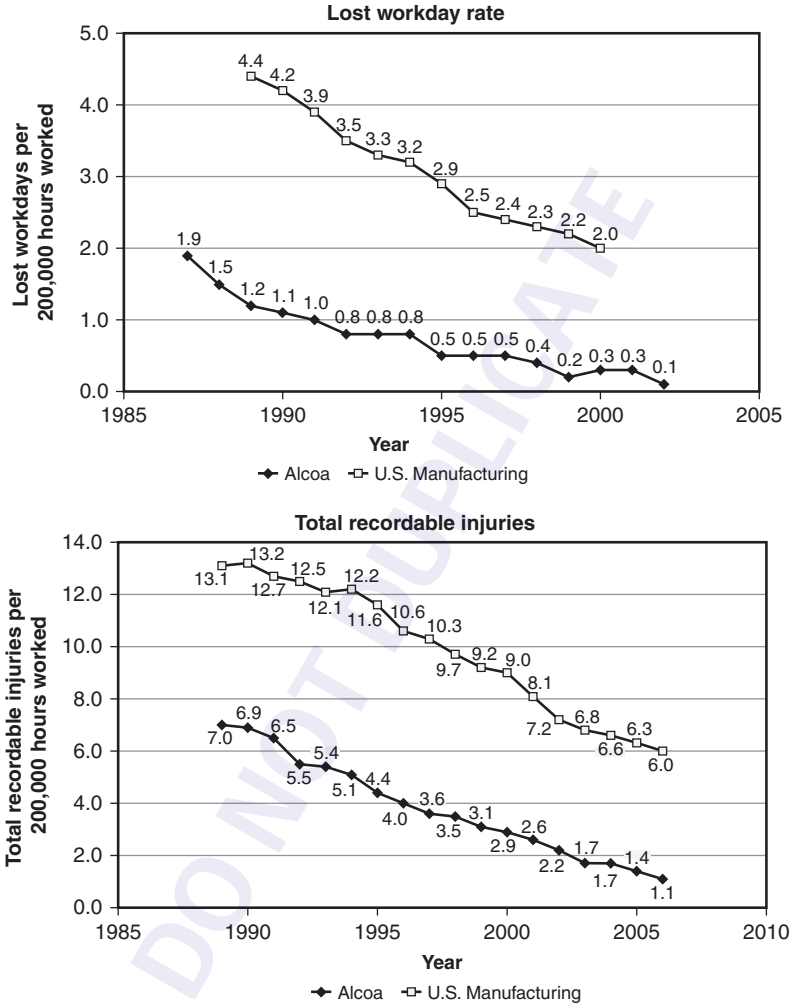
Producing aluminum products—soda cans, window and door frames, automobile wheel rims, and aircraft landing gear—requires that Alcoa use processes that would appear to be people-eating. Work begins in the bauxite mines with huge digging machines. Then the bauxite has to be refined into an intermediate product, alumina. This compound of aluminum and oxygen is not a usable commodity. It becomes valuable when it is dumped into containers called pots that are the volume of a railway car. Electrodes as big as telephone poles are jammed into the pots, delivering current that strips off the oxygen and leaves behind molten aluminum. With many scores of pots in a facility, the electricity used is enough to power a small city. But no one is in the market for liquid aluminum, so it has to be tapped and run into molds. Then it is reheated and stamped, forged, molded, rolled, or extruded under great pressure.

This combination of volume, mass, velocity, temperature, pressure, voltage, and current, with some caustic chemicals thrown into the mix, sounds dangerous. And at most companies engaged in such lines of work, it would be. Yet, Alcoa somehow defies those conditions. It is the safest large manufacturing employer in the United States, with a risk of an on-the-job injury that is one-twentieth of the national rate.

The graph in Figure 4-1 top shows the rate of lost workdays for Alcoa and for the overall U.S. manufacturing economy. This is the measure of a worker's chance in a particular year of getting hurt on the job seriously enough that he or she has to miss a day or more of work as a result. Even in the late 1980s, Alcoa already had an enviable safety record compared to the



Figure 4-1 Workplace safety at Alcoa: From safe to safest



nation as a whole. But what is astounding is how much it outpaced the pack in the ensuing 20 years. Whereas the United States overall had a 60 percent cut in risk from 4.4 percent to



2 percent, Alcoa's reduction in risk was more than 95 percent, from 1.9 percent to less than 0.1 percent.

On the more comprehensive measure of total recordable injuries, which includes less severe events that do not cause the loss of a day of work, Alcoa reduced risk by more than 80 percent, in comparison to a cut of 50 percent for manufacturing overall, as can be seen in Figure 4-1 bottom. And there is something else to keep in mind. Alcoa's progress during this period was not a trade-off, optimizing workplace safety at the expense of other measures. During the same period, Alcoa handily outpaced the Dow Jones Industrial Average (DJIA), with a stock price appreciation of nearly 700 percent, compared to approximately 470 percent for the Dow. (That Alcoa is a component of the DJIA, and so pulled up the average, indicates an even wider gap between itself and its large-market-cap peers). Alcoa did equally well when compared to a broader market index, the Standard and Poor's 500. Let's take a closer look at how Alcoa managed to tie exceptional improvements in workplace safety with outstanding economic performance leading to great market returns.

Back in 1987, the odds of getting hurt seriously enough to miss work at Alcoa were 2 percent per year. How bad was that? That meant that the odds of getting hurt in a decade were nearly 20 percent and that if you were going to make a career at Alcoa, the risk of getting hurt at least once would have been 40 percent over 25 years. With 90,000 workers at Alcoa at the time, it meant that seven or more workers were getting hurt on the job every day, approximately one per business unit. That was a hard responsibility to bear, particularly in a company in which it was not uncommon for neighbors and family members to work together.



Historically, there had been a view within the company that processes involving such complex chemistry and physics are inherently unstable and unavoidably dangerous. Certainly, you've encountered versions of this attitude in many other sectors, where the particular product, process, markets, and people—including employees, customers, and patients—are blamed for compromises in quality, safety, effectiveness, efficiency, and responsiveness that actually result from the failure of leaders to manage complex work systems for high performance.

All the same, there was a growing discomfort with the rate at which colleagues, friends, neighbors, and family members were being hurt. Alcoa's began to reexamine their assumptions. Perhaps harm was not inevitable. But in that case, what were the causes?

The idea that the processes were basically safe but that workers were deliberately self-destructive was rejected. So was the hypothesis that the workers were not smart enough to work safely. The record suggested that people got hurt not because they were stupid but because they found themselves in circumstances in which it was easy to get hurt and hard to be safe. (Remember Mrs. Grant's nurse in Chapter 3?) If the workers were not at fault, perhaps it was the research scientists and design engineers. Could they have designed safer processes? But no one believed they had deliberately failed to do so. The only explanation left was that Alcoa's processes and work sites presented unacceptable levels of risk because the company's scientists and engineers did not know how to design processes and workplaces correctly and its supervisors and operators did not know how to run them well enough.



This was a huge mind-shift. Like AT&T with Bell Labs, Xerox with its Palo Alto Research Center (Xerox PARC), and IBM with its research center, Alcoa was an industrial giant with a deep commitment to cutting-edge research and development. For years, Alcoa had been hiring top doctoral candidates in materials science, engineering, and industrial engineering from top universities and training them at Alcoa's Technical Center. If those geniuses did not know how to design a safe system, who did?

Alcoa was now on the verge of understanding one of the cornerstones of managing complex operations for high performance: No team can *design* a perfect system in advance, planning for every contingency and nuance. However, as Alcoa realized, people can *discover* great systems and keep discovering how to make them better.

When Alcoans got hurt or had close calls, leaping away to dodge a splatter of molten metal in a smelting plant or ducking at the last moment to avoid being hit by a swinging boom, they did so because they found themselves in situations no one had anticipated during design, which had been done at a time and place far from the actual work. Idiosyncratic confluences and coincidences of people, processes, products, places, and circumstances could create a hazardous situation where none had been known to exist. This was a seminal insight.

The problem was not bad motives, incompetence, or anything of that sort. Rather, it was a lack of foresight rooted in the inherent impossibility of anticipating the myriad interactions among the components that make up complex systems of work. Despite all the effort put into up-front design, something will always be overlooked. If it is impossible to be completely knowledgeable, ignorance is inevitable. However, it is



not irreversible. At Alcoa, people came to realize that behind ignorance lay opportunity. If Alcoans could spot unanticipated situations when and where they occurred, they could bring to bear the same disciplined knowledge-building behavior they exhibited in the R&D labs and get better processes as a result. The key was to identify problems as they occurred—the more, the better—and solve them when they were seen. If you had to depend on a single explanation for Alcoa’s success, it would be that Alcoa gave up depending on *designing* perfect processes and committed itself to *discovering* them instead.

## The Four Capabilities at Alcoa

This idea of seeing problems and then solving them was operationalized in myriad ways, none of which should be held out as a universally correct or comprehensive method. From what I’ve written earlier in this book, you know that I’m critical of those who try to achieve great outcomes by copying the specific solutions other people have developed for their own idiosyncratic problems. Look instead at the reasons why these solutions were successful where and when they were used. A good way to do that is to see how Alcoa’s policies and actions helped it to develop and use the four capabilities necessary for high-velocity management of complex organizations.

### Capability 1: Seeing Problems as They Occur

In 1987, Alcoa announced the hiring of a new CEO, Paul O’Neill. From the start, his approach was unusual. One might expect a new



CEO to establish corporate goals involving stock price, market share, return on assets, or return on investment—financial measures of success. Not O’Neill. As he announced in his first public appearance to the media and “the Street,” his primary concern at Alcoa would be safety. What might be a reasonable safety goal for a company engaged in so many perilous processes? How about reducing injuries by half? How about moving Alcoa into the top quartile, decile, or percentile compared with its peers? O’Neill ignored such relative measures. The goal was to be zero injuries to employees, contractors, and visitors. Why zero? Zero injuries meant perfect processes based on perfect knowledge of how to do work. Anything less than zero meant imperfect processes, and imperfect processes implied imperfect knowledge or ignorance. Therefore, when ignorance was found, it had to be rectified.

O’Neill and his colleagues built their strategy from this fundamental realization that things go wrong because there is insufficient understanding of how to make them right. As one way of acting on this belief, they insisted that within 24 hours of someone getting hurt in an Alcoa facility, something that was happening up to seven times a day, O’Neill had to be notified. (Over time, the reporting threshold became lower, including not only injuries but close calls or any unexplained worsening of someone’s condition that caused him or her to miss work.) However, it was not just that O’Neill was a data geek eager to track trends and tendencies or a megalomaniacal control freak dying to look over everyone else’s shoulder. He wanted to know within 24 hours because of the dynamic it would establish within the organization.

The kicker was that the reports had to come directly from the business-unit presidents. Why? After all, Alcoa had oper-





ations around the world; O'Neill might be anywhere at any given moment. This was a deliberate effort to create urgency around seeing and solving problems. For a business-unit president to inform O'Neill within 24 hours, he or she had to know about the problem well in advance of that deadline. This meant that the president had to hear from the vice president within an even shorter time frame, and the VPs had to know about injuries from their direct reports quickly enough to reach their unit presidents. When you consider the number of layers in the Alcoa hierarchy, this means that the first-level supervisors had to turn to the frontline employees and insist, figuratively if not literally: "If news of your injury is to make it to O'Neill in a day, you had better start yelling the moment you get hurt, before the pain sets in, maybe even before you are sure you have been injured."

What was this all about? O'Neill's 24-hour policy not only conveyed urgency but also encouraged accuracy. The sooner a problem is flagged, the more "perishable" information can be collected about it. Remember our reflection in Chapter 3 that if the staff at Mrs. Grant's hospital had waited to take stock of what had happened, empty vials would have been disposed of, memories would have faded, and they might never have reconstructed what had doomed the patient? In an industrial process, there is also the issue of information perishability. Temperature may change, pressure may drift, voltage or current may ebb or flow, and speeds may pick up or slow down. Enough drift and change, and the situation may be so different at the time of investigation that it is impossible to re-create the conditions associated with the failure and thus impossible to determine the cause. Without a known cause, treatments will



be arbitrary at best and recurrences will be likely. Even if machine conditions remain unchanged, people's memories are faulty and fade rapidly. In other words, if you do not see a problem when and where it occurs and swarm it to investigate it, much of the information needed to understand it will perish, spoil, fade, and dissipate. Once that happens, it becomes impossible to re-create the problem, nail down what caused it, and take corrective measures that will prevent its recurrence.

## Capability 2: Swarming and Solving Problems As They Are Seen

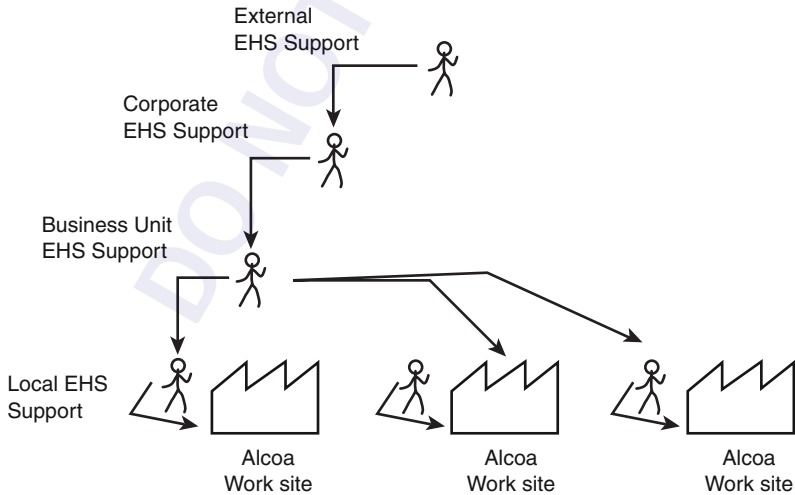
For the reasons just stated, there was a second rule: Not only were the business-unit presidents required to inform the CEO of an injury or near miss within a day, but within two days they had to report what the initial investigation had revealed about its causes and what was being done to prevent the problem from recurring. When a code team in a hospital races to an ill patient, they quickly size up the symptoms, immediately begin a diagnosis to determine what caused the symptoms, begin a treatment based on the diagnosis, and begin monitoring its effectiveness. To wait would risk misunderstanding the situation and leaving it uncorrected for too long. Alcoans learned to go through a similar, disciplined cycle of real-time problem recognition, diagnosis (*root-cause analysis* in industrial parlance), and treatment (*counter-measures* or *corrective measures* in manufacturing vernacular). It was the discipline of the Shewhart cycle—plan, do, check, act—popularized by Edwards Deming, but accelerated to warp speed.

This emphasis on rapid identification and swift investigation of safety-related problems was backed up with a commitment of



skilled resources. For all the technical expertise Alcoa had in the processes they designed and operated, many lacked the complementary knowledge of how to develop a safe work environment and foster safe work behaviors. Therefore, Alcoa invested in developing multiple layers of environmental, health, and safety (EHS) expertise that would be available when and where they were needed. If there was an injury or a near-miss in a facility, the shop floor workers and production engineers could get assistance from on-site experts. If that expertise proved insufficient, there was a pool of experts at the facility and business-unit levels who could pitch in. If they could not crack the case, Alcoa’s corporate staff would dispatch additional support, and if that proved insufficient, outside experts would be contracted to the team, as we see diagrammed in Figure 4-2. The key was to main-

**Figure 4-2 Environment, health, and safety expertise in support of “see every problem, solve every problem”**



tain the urgency to see problems, swarm them when seen, solve them when swarmed, and—as we will see with Capability 3—quickly spread the new knowledge throughout the organization.

### **Capability 3: Spreading New Knowledge**

This high-velocity approach of seeing problems and solving them when and where they occur proved pivotal for Alcoa. No longer burdened by the attitude that things inevitably go wrong when people work with large-scale industrial processes, Alcoans gradually stopped working around the difficulties, inconveniences, and impediments they experienced. Coping, firefighting, and making do were gradually replaced throughout the organization by a dynamic of identifying opportunities for process and product improvement. As those opportunities were identified and the problems were investigated, the pockets of ignorance that they reflected were converted into nuggets of knowledge. That knowledge had a special quality that was of great competitive significance.

Alcoa was hard-pressed to distinguish itself from its competitors by positioning itself uniquely relative to its external environment. Exclusionary contracts for bauxite were not an option. Electrical power and the chemicals used in refining and smelting were commoditized and the basic processes of making aluminum had been known for decades. Alcoa was subject to the same regulations as its competitors. Certainly, customers did not want to be drawn into a monopolistic dependency on Alcoa.

However, by seeing problems and solving them in an accelerated fashion, Alcoa was building process knowledge that was not only hard won, but also scarce and proprietary—unavail-



able to outsiders who did not make the same efforts. Since the more use Alcoa could make of these discoveries, the more valuable it would be, Alcoa made sure that what was discovered locally was shared organizationally.

There were many mechanisms for this. First, of course, were the many cross-fertilizing “honeybees” Alcoa created by emphasizing the rapid identification, reporting, investigation, and resolution of safety-related problems. As new problems were sped up the managerial ranks, they came to the awareness of people who might have seen something similar in another part of the company for which they were responsible. Therefore, they could lend help, assistance, and insight, spreading knowledge from one area to another. Certainly the environment, health, and safety experts helped this pollination process, carrying the lessons they learned in one area to another.

Then there were the deliberate attempts to ensure that what was learned locally had benefit systemically. Just as Alcoa defied convention when it established safety, rather than a more traditional financial measure, as its top priority, it did so again when it instituted its first corporate-wide information technology system. Unlike companies that might have made accounting, payroll, taxes, benefits, or another financial function the first corporate problem to be solved with IT, Alcoa tackled safety first. The idea was that no matter where you were in Alcoa, if you had an incident, you could make it visible to anyone else in the company and if you had a problem, you could investigate what others who had had a similar experience had learned. When people did their work at Alcoa, they were drawing on much more than their individual expertise.



Individual performance could reflect the collective experience of the organization.

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

In most organizations, middle managers play an essential but bureaucratic role. They convey high-level goals that are set at more senior ranks, restating them as objectives relevant to the part of the firm for which they are responsible. From the lower ranks they convey information upward, taking specific data and reformulating those data so that they can be used by corporate decision makers. These middle-management roles make it possible for sprawling organizations to allocate resources and coordinate activities, as has been documented by Alfred Chandler and other business historians. However, Alcoa was not content to let middle managers be information conduits and coordinators, nor was it content with the model of scientific management championed by Frederick Winslow Taylor in which the “brains” of the organization developed optimal procedures for the “brawns” to employ.

Instead, Alcoa expected its leaders at all levels to develop the organization’s ability to manage work in such a way as to see problems, solve problems where they were seen in order to build new knowledge, and spread that knowledge so it would be useful throughout the organization. Leaders not only had to have detailed process knowledge in their own right, in order to understand what was occurring and why, but they also had to coach and train others to be able to see deficiencies in how work was conducted and then develop and validate cor-



rective measures. In this formulation, managers not only oversaw the production of physical goods with capital equipment, without doubt an essential role, but they were also responsible for cultivating the skills of inquiry and invention necessary for generating intangible assets—the process knowledge that would set Alcoa apart.

This approach was not merely preached; it was practiced. Careers rose and fell in accordance with how well the more senior managers could develop the capabilities of their direct reports and deploy those capabilities to increase safety and efficiency and reduce environmental impacts. In fact, in one dramatic case, a business-unit president was dismissed over safety-related issues.

The unit he ran had been very successful by most measures. Revenue had grown dramatically and customers were delighted. What cost this otherwise successful leader his job was that, on his watch, a worker in an assembly plant had gone home early, feeling nauseated. The cause of the nausea was unexplained and the man returned to work the next day, apparently unaffected. Two weeks later, several more workers went home early suffering from nausea. The investigation that followed revealed that an idiosyncratic set of circumstances had resulted in indoor air pollution that, in turn, had caused the workers' symptoms.

The business-unit president lost his job because when the first incident went unexplained, it went unreported, insufficiently investigated, and unresolved. Help was not pulled in to bolster the investigation. Other people at the same site could have been exposed to the same unknown risk; in fact, they were. The new business-unit president was chosen precisely



because he was deemed able to reinvigorate the “see a problem, solve a problem, share what you have learned” dynamic. Sure enough, when the dynamic was restarted, people began to discover latent problems.

## Sustaining and Expanding the Results

Alcoa moved from an approach in which problems are accepted as unavoidable—the “one thing after another” we expect with complex operations—to an approach in which problems are clear signals, beneficent warnings, the system saying, “There’s something important you don’t know about me, but if you listen, I’ll tell you.”

Over the course of 20 years, Alcoa cut its rate of on-the-job injuries leading to a lost work day from 2 percent to 0.07 percent. Whereas the 2 percent rate meant that senior managers learned every day that someone had been hurt, now it was days and weeks between reports. For the shop floor worker, a risk of 0.07 percent translated into a chance of injury of less than 1 percent in a decade and meant that over the course of a 25-year career, the chance of getting seriously hurt on the job was less than 2 percent. In contrast, for a non-Alcoa employee, the risk of a lost workday fell from 36 percent in a decade and 68 percent in a career to 18 percent in a decade and 40 percent in a career.

It might be easy to attribute Alcoa’s success in improving workplace safety during the stewardship of Paul O’Neill to a singular focus by a charismatic leader. However, that interpretation requires dismissing several factors. First, improvements in safety as measured by total reportable incidents and inci-





dents leading to lost workdays continued to decline after O'Neill's departure. Second, Alcoa's intense focus on workplace safety did not mean that organizational improvement was measured by safety metrics alone. Improvement in safety did not come at the expense of improvements in quality, yield, efficiency, and cost, all of which contributed to improved profitability and market capitalization.

This ability to improve across the board, rather than improving safety at the expense of something else important, depended on the fact that focusing on workplace safety had both moral and practical motivations. The moral rationale, as described above, grew from a basic discomfort with putting people in harm's way. The practical rationale was that if people lacked sufficient knowledge to design and operate processes perfectly from a safety perspective, they probably lacked the knowledge to design and operate them perfectly in terms of quality, efficiency, yield, and timeliness. Thus, safety (or lack of safety) opened a window into all the underlying factors that compromised Alcoa's performance in terms of the measures more typically of concern to large industrial companies. During the period when Alcoa focused on safety, it improved other dimensions of performance as well. Earlier we discussed stock market returns. Now let's look more closely at specific examples.

Alcoa's engineered-products plant in Cressona, Pennsylvania, increased productivity on two lines by 87 percent by redesigning work flows, improving equipment, and developing better work methods. Packing costs were reduced, delivery performance was increased, and injury risk was cut. It was not that this plant specifically "managed safety" or that it specifi-



cally managed quality, on-time performance, or efficiency. It managed the processes it was using and thereby improved their performance by numerous measures simultaneously.

In a similar fashion, Alcoa's Davenport, Iowa, plant addressed increased demand through better process design. As with Cressona, the plant didn't make a trade-off between one good thing and another; rather, it learned to extract more yield from all its efforts. Ten extra hours of rolling time per month were freed, which had \$500,000 in value; inventory was reduced by \$1 million; and 19 improvements in environment, health, and safety were carried out.

Alcoa's continued focus on improving its processes, wherever and whenever the opportunity arose (that is, whenever a problem or unexpected outcome occurred), made it high-performing across the board. Safety measures continued to improve even as 2005 was marked by record revenue levels and improvements in return on capital and by recognition as one of "the best-practice leaders in cutting their greenhouse gas emissions." The year 2006 brought even better results: revenues up 19 percent to a record \$30.4 billion, income from continuing operations up 72 percent, and return on capital up to 13.2 percent. It was also the twentieth consecutive year of improvement in safety measures, with a 96 percent reduction in the lost workday rate since 1987 and an 88 percent reduction in total recordable incidents.

All this resulted from deciding that problems were not a never-ending plague to be endured but a never-ending guide to improvement. In the next chapter, we'll see how a similar approach to high-velocity management led to great results in a complex, hazardous military situation.



---

## REFLECTIONS ON MANAGEMENT AND LEADERSHIP

When I first met Paul O'Neill, I had what I have come to recognize as a naive view. As the CEO of an enormous company, O'Neill must have extraordinary power, I thought, along with an extraordinary ability to shape events and direct people. The bases for that expectation were several. There is the business media's celebration of the individual "captain of industry" and "titan of commerce," the singular hero who introduces products, resuscitates companies, and otherwise has a profound impact. In this view, Chrysler was "saved" by Lee Iacocca, Jack Welch single-handedly drove General Electric to new heights, and so forth. One would think that Bill Gates single-handedly wrote all the code at Microsoft and designed and carried out all its strategic maneuvers. (I remember people saying that at the very least he reviewed every line of code.) Or perhaps that every element, feature, and nuance of the Apple iPhone and iPod was put in place by Steve Jobs. We celebrate celebrity and fuel the myth of the leader as supreme architect, engineer, and pilot. These notions are strengthened in business school courses that posit management as a set of chesslike strategic transactions and discuss complex systems as being amenable to sophisticated mathematical modeling and control.

That is not at all what leadership is like in a process/systems-intensive organization operating over the long haul. I came to appreciate that the leader of a large organization does have tremendous power, but



much of it is of a destructive nature. He or she can fire people, shut down facilities, divest product lines, and disengage with difficult customers and suppliers. However, constructive power is harder to muster because creation ultimately is a collaborative and coordinated effort. Collaboration and coordination are tricky because they require that those who participate have at least some degree of agreement about what they are trying to achieve, what approaches are acceptable and preferable, and what is off-limits. Short of that, as Howard Stevenson pointed out, you are depending on despotism to get things done. Clear definitions of desired outcomes and clarity of roles and methods are needed even in fairly small organizations such as orchestras, dance companies, and bands. How much more so in organizations that have hundreds if not thousands of people contributing to the achievement of a larger whole?

For someone in Paul O'Neill's position, anything he or she says will be repeated, but with imperfect duplication. By the time an executive pronouncement is repeated and relayed, it will be distorted and misframed. If the leader is trying to achieve something significant, the countermeasure to distortion is to "broadcast" the key message consistently and repeatedly so the "signal" will emerge from the static and noise that develop with each successive round of transmission. In Paul O'Neill's case, this meant that his consistent message, which was not going to compete with many others, had to be that safety was a primary concern, zero injuries was the goal, and the identification of and adherence to safe practice had to be exceptionally



rigorous, with immediate identification and resolution of threats to safety as the means to better outcomes.

Two images come to mind. The first is attending a kickoff meeting in Alcoa's former corporate headquarters for the Pittsburgh Regional Healthcare Initiative, of which O'Neill was a founder. Before starting the meeting, O'Neill stated that he saw many people in the room who had not been in the building previously and he wanted them to be assured of their safety. Thus, to an assembly of 40 to 50 people, the chairman of one of the world's most prominent companies explained where the exits were, what to do in case of an emergency, and how to leave the room, the floor, and the building safely.

The other image gets to leadership's impact on culture. I was in a large Alcoa extrusion plant in Brazil in which aluminum logs are forced under great pressure through a series of dies to create window and door frames. There is heavy material, heavy equipment, and loud machinery. In the middle of the tour with a number of senior-level executives, I had trouble hearing the guide's explanation. I pressed closer, but when that did not help, I pulled my earplugs out so that I could hear better. In only a few moments, one of the operators came over and in a combination of Portuguese, English, and creative pantomime, indicated that I had to wear hearing protection or leave the production area. I was struck by the cultural chutzpah that was necessary for that to occur.

If I ended the characterization of leadership at this point, it might seem that an effective leader is one who "manages by objectives," with a few other platitudes



thrown in. However, that is an overly reductionist view, making it sound as if the right set of objectives, repeated ad nauseam, will lead to great outcomes.

There is another piece to my picture of leadership: energetic attention to detail, grounded in the belief that leaders have to have a deep understanding of how things work if they are to develop people, guide people, and make decisions. For instance, in a talk before his retirement, O'Neill discussed the thermal inefficiencies of producing aluminum using current processes, the impact on cost, and the ability to sell aluminum for applications beyond those for which it currently is used. There was a logical thread between British thermal unit efficiency, the costs of production, industry capacity, and the effects on supply and demand.

We'll see in the next chapter how this commitment to managing from a few simple but robust principles, coupled with tremendous attention to detail and the development of people, is played out in other high-performing organizations.

---



## CHAPTER 5

# HIGH VELOCITY UNDER THE SEA, IN THE AIR, AND ON THE WEB

In Chapter 4, we looked at how one company, Alcoa, managed its complex systems of work to see problems, solve problems, and share what was learned, all the while insisting that leaders cultivate these capabilities. In this way, Alcoa accelerated the rate at which it learned how to design and operate its technical processes and systems of work, thereby achieving exceptional performance. And though it started by focusing on problems related to workplace safety, it soon found that safety problems reflected process ignorance and that this ignorance would also manifest itself in other problems such as quality, timeliness, and yield versus scrap.

In this chapter, we'll look at three other organizations that used the velocity with which they created and employed useful knowledge as the basis for achieving exceptional performance. The first case is the U.S. Navy's Nuclear Power Propulsion Program, which invented, introduced, and operated an exceptionally challenging technology with greater



speed and reliability than organizations charged with comparable challenges. In the second case, Pratt & Whitney accelerated its process for bringing new jet engine designs to market. The third case concerns a pioneer dot-com company which survived the 2000 market shake-up, established itself as a profitable enterprise, and wound up converting a small initial investment into a fortune. As widely as these examples differ in their missions and circumstances, they all illustrate how high-velocity organizations achieve superlative outcomes by applying the principles delineated in this book.

## **U.S. Navy Nuclear Power Propulsion Program**

The U.S. Navy has launched more than 200 atomic-powered ships—using up to 30 different power plant designs, with 500 reactor cores brought into operation—since the start of the nuclear power propulsion program in 1948. As of 2006, those ships collectively have had more than 5,700 reactor-years of operation and have “steamed” well over 134 million miles.

This in and of itself is a technological and managerial marvel considering what came before. In World Wars I and II, submarines were a strategic threat, sinking substantial merchant marine traffic and, by the fear they aroused, forcing military and commercial convoys to take extraordinary precautions on open-water voyages. Watching Hollywood renditions, one might conclude that those subs were lethal because they could remain hidden for extended periods, sneak up on their prey undetected, and attack with devastating force.





That was not the case; the Hollywood image overplays the capabilities of submarines and downplays their vulnerability.

In truth, the performance of submarines was limited by the batteries that powered them when they were submerged. The batteries held charges only for short periods, so underwater range was no more than 20 miles. Most of the time, the subs were forced to operate on the surface, where they had air to run their diesel engines but were exposed to detection and destruction by larger warships and aircraft. In real life, success often meant sneaking in close, remaining submerged only briefly, then compensating for the ineffectiveness of the torpedoes by fighting a close-in battle with small mortars and machine guns mounted on the decks.

Under fire or not, life for submariners was difficult, even by the Spartan standards of military craft, which are largely designed to move weapons systems with maximum effectiveness, only accommodating the crew as best as they can. Submarines, being smaller than other ships and designed for underwater operation, ran rough on the surface. Once they were under way, the demands of power conservation meant poor ventilation and often moldy food. Their cruising range was limited by the amount of fuel they could carry; before it was gone, they had to stop to refuel at sea or in port.

Nuclear-power propulsion erased those limitations. Nuclear-powered submarines have scored repeated milestones: submerging below the polar ice cap, traveling beneath it from the Pacific to the Atlantic, rendezvousing with other submarines underneath it, surfacing through it, and circumnavigating the Earth completely submerged. Whether used for intelligence gathering during the Cold War, deployment



of special forces, tracking of Warsaw Pact warships, controlling sea-traffic choke points during a conflict, or carrying intercontinental ballistic missiles and thus guaranteeing a retaliatory strike capability, nuclear-powered submarines changed the fundamentals of naval doctrine in the post-World War II era thanks to their ability to remain submerged almost indefinitely.

Nuclear power also revolutionized aircraft carriers, which had earlier revolutionized ocean warfare during World War II. The dreadnought battleships of all navies fell before the onslaught of seaborne air forces, which could project force farther and faster. In the battle of Midway, one of the decisive sea battles of the Pacific, the opposing fleets never fired on each other directly; instead, the aircraft of each one attacked the ships of the other. Carriers could also provide air cover to soldiers and marines “storming the beach” before airfields could be secured that were within flying range of the conflict’s leading edge. This helped offset the advantages land-based defenders gained with artillery and their own air power.

Putting nuclear power aboard aircraft carriers was another order-of-magnitude change in the strategic balance. It provided the additional advantages of range, speed, time on station, and ability to conduct unlimited launchings and landings. The U.S. Navy’s ability to police sea-lanes, keeping them open for commerce, and to project military power when and where necessary was greatly enhanced.

If we stopped here, the introduction of reliable nuclear propulsion onboard warships would be a remarkable accomplishment in its own right. Yet there are other considerations that should draw our attention. The first is the extraordinary



velocity with which this technology was introduced. The first nuclear-powered submarine, the *Nautilus*, entered the fleet in 1954, a mere blink of the technological eye when one considers that it was not known how to harness the atom only a decade previously and that the program to develop nuclear propulsion did not take form until 1949. It is all the more remarkable when one considers that launching this brand-new technology required the discovery of new science, the invention of new materials for shielding and reaction control, the creation of new manufacturing systems, the design of novel devices and power plants, and the training of thousands of engineers, craftspersons, and operators. The technological and organizational accomplishments were fantastic.

And yet there is an added wrinkle. Since the launch of the *Nautilus*, the Navy hasn't suffered a single reactor-related casualty or escape of radiation—a far cry from what comparable programs have experienced. Before the dissolution of the USSR, the Soviet fleet suffered a number of nuclear calamities in its 50-year history, with substantial injury, death, environmental pollution, and destruction of equipment (see Table 5-1 for examples through the 1980s). NASA, undertaking a comparably difficult and dangerous mission, has lost one Apollo crew and two shuttle crews in a little under 50 years of manned space flight. As we saw in Chapter 3, the civilian nuclear-power industry has hardly been trouble-free.

How can this extraordinary performance be explained? Attention naturally turns to the demanding and monomaniacal commitment of the founder and longtime leader of the Navy's Nuclear Power Propulsion Program (often referred to as "NR" for *Naval Reactors*). Hyman Rickover, a 1922 graduate of the



**Table 5-1 Partial List of Calamities in the Soviet Nuclear Navy**

<i>Ship</i>	<i>Date</i>	<i>Problem</i>	<i>Consequence</i>
K-8	Oct. 13, 1960	A leak developed in the steam generators and in a pipe. Equipment for blocking leaks was damaged. The crew began the work of stopping the leak.	Large amounts of radioactive gases leaked out, contaminating the entire vessel. Three of the crew suffered visible radiation injuries.
K-19	July 4, 1961	A leak developed in the primary cooling circuit, causing a drop in pressure and setting off the reactor emergency system.	The crew worked long periods in radioactive areas of reactor compartment. All were exposed to substantial radiation. Eight died.
K-11	Feb. 1965	While the submarine lay in dock, the reactor lid was lifted without control rods being secured. While the problem was being investigated, it happened again.	There were releases of steam, and a fire broke out. The reactor was retired and replaced.
K-27	May 24, 1968	Power inexplicably dropped suddenly during sea trials.	Radioactive gases were released, and radiation onboard increased. The reactor was shut down and approximately 20 percent of the fuel assemblies were damaged. The ship was scuttled in 1981.
K-140	Aug. 1968	Wrong installation of the control rod cables and error.	Unplanned automatic reactor start-up while in shipyard.



K-429	1970	Uncontrolled start-up of ship's reactor while submarine was at shipbuilding yard.	Fire and release of radioactivity.
Echo-I sub.	Aug. 21, 1980	Vessel suffered a radioactivity leak following a fire.	Nine crew members died and three others were injured.
K-222	Sept. 30, 1980	Breach in procedure let power through safety-rod mechanism without the controls being engaged. Automatic-equipment failure caused uncontrolled raising of control rods and uncontrolled reactor start.	The reactor core was damaged.
K-123	Aug. 8, 1982	Leak in steam generator caused release of liquid-metal coolant into reactor compartment.	Reactor had to be replaced. It took nine years to repair the submarine.
K-314	Aug. 10, 1985	Control rods incorrectly removed when the reactor lid was raised. Reactor went critical during refueling.	Explosion released large amounts of radioactivity. Ten people were killed.
K-431	Dec. 1985	The reactor overheated while the vessel was returning to base outside Vladivostok.	Submarine is now laid up at the naval base in Pavlovsk.
K-192	June 25, 1989	A leak was discovered in the primary circuit; the submarine surfaced. Botched attempts to fix the leak and cool the reactor led to a cascade of misfortunes.	Releases of radioactive iodine were detected in the areas surrounding the submarine and later on land.



Naval Academy, did not retire from the Navy until 1982, making him the longest-serving officer in the U.S. Navy's history. He created for himself an exceptional position of autonomy and power with two appointments: a civilian one from the Atomic Energy Commission (AEC), which was responsible for the design, development, and deployment of nuclear power generally; and a military one within the Navy, which was responsible for the contracting, design, construction, and operation of warships. In effect, he put himself in a position to make demands on the Navy from his AEC perch and on the AEC from his Navy perch. He cultivated relationships with members of Congress responsible for allocations and promotions and had influence on the budgetary process that often outweighed that of his civilian and uniformed superiors. He had influence over defense contractors as well, given his administrative power over major research, design, construction, and maintenance programs.

Furthermore, Rickover was intimidating. His interviews with prospective members of the reactor program and potential officers on nuclear-powered ships were legendary. One story is that he cut an inch from the front legs of the chair in which interviewees sat to make them feel uncomfortable without knowing why. He was known to berate and insult candidates for the program.

Theodore Rockwell, who was part of the initial group that started the nuclear-propulsion program with Rickover, wrote in his memoirs about a call between the two of them, conducted over the single undersea phone cable. Rickover grew increasingly frustrated with the poor sound quality, screaming at Rockwell until the operator finally cut in, "Sir, if you would just speak in your normal voice. . . ."



“Goddammit, this is my normal voice!” screeched Rickover. To which Rockwell added: “That’s true, operator, it is.”

Finally, there are any number of accounts (some of which we will encounter later in this chapter) of Rickover’s determination to know every detail of the technology for which he was responsible, including whatever mishaps befell it. He certainly can be perceived as an archetypical unpleasant, condescending, micromanaging boss.

But however hard-driving, cantankerous, or brilliant Rickover might have been, he cannot be the entire explanation for NR’s success, simply because he could not have solved every problem—at least not every technical problem—on his own. Furthermore, in the 26 years since his retirement, there have been several successors and countless civilians and sailors have served in NR without having known him first-hand—a good number wouldn’t even have been born before he left—yet the program’s perfect safety record has been maintained.

It can’t just be Rickover; there has to be something about the way in which the nuclear program’s complex work was and still is managed. And so there is. In response to the outrageous challenge it faced, NR developed what one of its chroniclers, Francis Duncan, called “the discipline of engineering.” This discipline was required because, whatever knowledge the group had, it was assumed to be inadequate. Therefore, there was no room for guessing; learning had to be constant and fast, not only experiential but experimental. To accomplish this, NR had to make explicit its best understanding and expectation of what actions would lead to what outcomes. Ensuring that people started with the best possible knowledge built into their



approach increased their likelihood of being successful. It also increased their opportunity to learn. With expectations clear, it would be obvious when something happened that didn't conform to those expectations. As a result, even if you didn't succeed, you created an opportunity to learn to succeed. Stating clear expectations was a given, with no exceptions; that's what made it part of the *discipline* of engineering.

With expectations clear, NR had to identify immediately when its best understanding was faulty—another discipline. And with equal discipline, each clearly identified pocket of ignorance was to be converted into usable knowledge. Finally, that knowledge had to be incorporated into updated designs for machines and procedures throughout the fleet; this, too, had to be done with rigor and discipline. All the while, this discipline of engineering was to be modeled, taught, encouraged, and harnessed by both junior and senior leaders.

When all this happens consistently, it changes the basic dynamics of an organization. Rather than letting each experience be either a success or a failure—but in neither case improving anyone's chance of success on the next try (see Figure 5-1)—every experience is *designed* to increase the likelihood of success on the next try as knowledge and know-how accumulate (see Figure 5-2). This was Alcoa's approach, it was NR's approach, and, as we'll see later, it is the consistent approach of high-velocity organizations more generally.

In Chapter 4, we saw how Alcoa's practices mapped onto the four capabilities first mentioned in Chapter 1. We'll now see how the practices Rickover instilled in the Naval Reactor program also did.





Figure 5-1 Succeed or fail


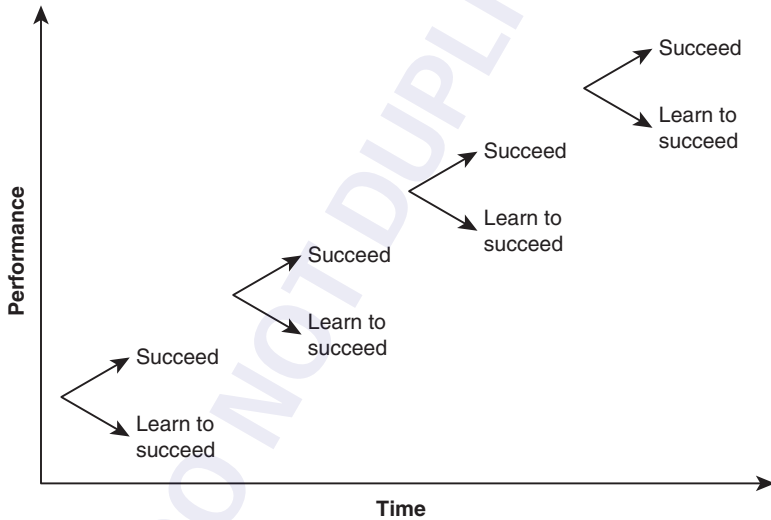


Figure 5-2 Succeed or learn to succeed



## Capability 1: Capturing the Best Collective Knowledge and Making Problems Visible

What do the terms *incident* and *incident report* bring to mind? An accident, an injury, a fatality, or damage to prop-



erty? Now ask yourself, why would an organization insist on incident reports? Accountability, reprimand, punishment? If your answers were any of the above, you are far from the Naval Reactor program's approach. It has a much lower threshold of what it deems an "incident" and a much higher threshold of what must be done when an incident occurs.

Let's start with a simple example: people working in a system that had already been designed and built and was now in service. Those who operated reactors onboard a ship and those who conducted maintenance and refueling onshore were expected to follow scripted procedures with exacting accuracy. There were clear expectations about what each person and each piece of machinery would do, in what order, and with what effects on each other. Incidents were strictly defined as departures from procedure. If they occurred, they had to be reported. For example, if someone were to start step 3 before receiving the agreed signal that step 2 is done, that would be an incident. And however inconsequential the outcome, an incident had to be reported.

This wasn't simply bureaucratic housekeeping. Just as close calls at Alcoa were indications that something about a manufacturing process was not completely understood, an incident in the nuclear navy meant that something about the way work was done was incompletely or inaccurately understood. This ignorance could not be tolerated. That part of the system could be connected to other pieces in ways that were not well understood; an incident that seemed inconsequential in isolation might be disastrous in just the right combination with other incidents.



This discipline of specifying expectations was not just for the frontline work of operating the submarines. It applied to everything and Rickover himself modeled this way of life.

Rockwell recalls preparing for a meeting and being challenged by Rickover to describe how the meeting would conclude—before it had even started. Rickover was not hazing him and did not expect him to be clairvoyant. Rather, he wanted Rockwell to predict in advance what a successful outcome would look like and how he expected to get there so he could determine whether something was amiss as the meeting proceeded. What was there about the situation, the discussion, the technical content, or the discussants that he had misunderstood? What were the consequences of that misunderstanding? What had to be done to address those misunderstandings? Those were all critical concerns, which otherwise might have been missed had Rockwell not been prepared to be surprised by events unfolding contrary to what he had anticipated.

Even—and especially—in upfront design and development work, where there were obviously great gaps in what was known about a particularly complex situation, this discipline was required. Rockwell describes designing the radiation shielding for reactors (a topic on which he became expert enough to author several books). No one knew how neutron bombardment would fatigue the metal and how the piping's welds, joints, and bends would affect radiation patterns. Therefore, when it was time to test the shielding, a grid was laid over the surface, with sensors distributed all across it. But the evaluation didn't rest at that.

Before any measurements were taken, Rockwell insisted that predictions be made about what the measurements at each



point would be. It was not sufficient to find out if the various sections passed or failed in terms of emitted radiation. Rockwell and his colleagues already knew that they would be wrong at many points since the science and technology were still in early stages. Therefore, they wanted to know for certain—sooner rather than later—exactly where and when they were wrong and what they misunderstood. The sensors were not just there to mark safe and unsafe situations. They were there to identify pockets of ignorance on the part of the shielding designers.

That is why, rather than just recording readings and noting where the exposure was too high, they first predicted what the readings would be and then compared those predictions to the actual readings to discover where their understanding was confirmed and where it was refuted. If the shielding worked less well than needed or expected, that certainly warranted investigation and additional engineering. We would all recognize that. However, if the shielding worked better than needed or expected, that, too, revealed a gap in their knowledge which could prove costly or dangerous and which needed to be plugged. It is not clear we would all see that as a learning necessity as well. The difference? Many tests are meant to distinguish good from bad. In this case, Rockwell structured the test to distinguish *understood* from *not understood*.

Similarly, when it was unclear whether hafnium or a silver-cadmium alloy would be preferable for controlling the rate of chain reactions, the choice was driven by comparative trials. But those trials were not simply tests to see which material was better than the other. Before the trials were started, the engineers predicted how each would perform, explaining why they thought so. The point was not just to make a choice between one material and the



other, but to identify things still not understood about both along the way. (In Chapter 7, we will see the same point made.)

From its very beginning, the Naval Reactor program was committed to following the script it had created without exception because making an exception would be to knowingly back away from the best understanding that had so far been acquired and would needlessly confound any analysis of future experiences. Consider what happened in September 1954, when *Nautilus* was being tested only a few months before its launch date. A steam pipe burst. The investigation showed that the wrong kind of piping had been installed. The NR program had all of that type of pipe ripped out and replaced with what had been specified. There was no thought of testing to see if some of the “wrong” pipe might pass some arbitrary performance test anyway. Until there was further disciplined study, NR couldn’t be confident that the other pipe was adequate to the demands that would be placed on it.

NR’s reaction did not stop with containment—replacing the wrong pipe with the specified pipe. That would have been a workaround that would not have targeted the underlying factors that were at the root of the problem. How had this mistake been made? What was to stop it from happening again? What was to stop it from happening with some other material or component? The progression of whys and hows traveled back through the value stream and supply chain. Until NR could answer all these questions, there was a deadly booby trap somewhere in its operations.

Not only did NR demand a high degree of specification of what was thought to lead to success, it wanted to be sure that when something was amiss, that too was clear. High-velocity



organizations are not in denial about human imperfection. What they want is operations that will snitch shamelessly—not only loudly, but accurately and quickly.

We see this in a report prepared by NASA, the background of which has a certain painful irony. More than 15 years after the 1986 *Challenger* disaster, NASA embarked on a series of benchmarking studies to understand how other organizations had achieved extraordinary levels of safety despite the hazards of their work. Subjects included Alcoa and Bath Iron Works, on the topic of workplace safety, and the Navy's nuclear-power propulsion program, its software-integration program, and the SUBSAFE program referred to later in this chapter. Between the first few studies and the last, there is a several-year gap when the *Columbia* tragedy interrupted the benchmarking effort. One cannot help but wonder if that disaster would have been averted had NASA started its studies earlier.

This is what NASA observed when comparing the Navy's design of nuclear reactors with the civilian approach at Three Mile Island:

In the case of Three Mile Island (TMI) commercial reactor, over 50 alarms or warnings were active prior to the mishap. At the onset of the TMI event, 100 more alarms were activated (a total of 150 of about 800 alarms active). In contrast, the total number of alarms and warnings in an NR reactor system is strictly limited to those needing an operator response. The Commanding Officer must be informed of unanticipated alarms that cannot be cleared. Naval nuclear power plants do not routinely operate with uncorrected alarms or warnings.



At Three Mile Island, the system spoke up, but the staff had learned to work around chatter which could not be understood. NASA found that within the Navy, alarms are simplified, so they don't sound so often. But when they do sound, or when other things go wrong, they are taken seriously. The NASA benchmarking team observed the sheer frequency with which NR spotted and reported on problems as follows:

This system is thorough, requiring deviations from normal operating conditions to be reported, including any deviation from expected performance of systems, equipment, or personnel. Even administrative or training problems can result in a report and provide learning opportunities for those in the program.

NASA noted that NR had established an exceptionally low threshold for what counted as a problem or incident, as was mentioned before:

During a General Accounting Office (GAO) review of the NR program in 1991, the GAO team reviewed over 1,700 of these reports out of a total of 12,000 generated from the beginning of operation of the nine land-based prototype reactors that NR has operated.

And that 12,000 doesn't even include the far more numerous ship-based reactors. The NASA report continues:

The GAO found that *the events were typically insignificant* [emphasis added], thoroughly reviewed, and



critiqued. For example, several reports noted blown electrical fuses, personnel errors, and loose wire connections. Several reports consisted of personnel procedural mistakes that occurred during training activities. . . .

## **Capability 2: Building Knowledge by Swarming and Solving Problems**

In high-velocity organizations, the response to problems is frequent, serious, and disciplined.

In many organizations, such emphasis on reports and written documentation as described above might be dismissed as bureaucratic obsessive-compulsive command and control, particularly if reports were required but simply filed and ignored. This is not the case in the Naval Reactor program. NASA observed not only the frequency but the seriousness of these reports.

NR requires that events of even lower significance be evaluated. Thus, many occurrences that do not merit a formal report to headquarters are still critiqued and result in identification of corrective action. These critiques are reviewed subsequently by the Nuclear Propulsion Examining Board and by NR during examinations and audits. This is part of a key process to determine the health of the program's self-assessment capability.

This was not just paperwork and it was not delegated as grunt work to junior officers and enlisted personnel; it was treated as an essential part of leading others. When a ship was being evaluated and had done well in some but not all cate-





gories, it was the responsibility of the commanding officer to explain why in a special letter that listed the failures and detailed the corrective actions. Promises to do better or try harder were never enough. The NASA benchmarking team pointed out that during Rickover's tenure and that of his several successors, each report "identifies the necessary action to prevent a recurrence." But such actions could not be precipitous. Solutions at NR had to be found through a process of disciplined discovery so that they could be trusted and safely propagated throughout the organization.

(Imagine a nurse in the hospital where Mrs. Grant died filing a report that he had *almost* mistaken a vial of insulin for a vial of heparin—no harm done, just a close call—then the hospital's chief of nursing reporting that incident to the hospital's CEO, along with her explanation of how such a thing could happen, where else such a thing could happen, and what had been done to make sure that these things did not happen again.)

Leaders in charge of the Navy's nuclear-reactor program, like their counterparts at Alcoa, discerned that the only way to understand and improve what was poorly understood and in need of improvement was to create ample learning experiences—the more the better, the sooner the better, the faster the better, the cheaper the better, and the greater the clarity of cause and effect the better. This, too, was done with great discipline. A fundamental aspect of the approach was not to take the obvious for granted but instead to make latent assumptions explicit and then test them. The opposite approach was used at NASA: The original assumption that foam shedding from the external fuel tank posed a threat to the heat shields was gradu-



ally replaced by an assumption that it did not. The second assumption may have been *buttressed* by the fact that nothing terrible had happened so far, but it had never been rigorously *tested*. Francis Duncan, the Rickover biographer cited earlier, made a point of how reality, not human assumptions, was always to be given the final word (since reality will always have the final word whether we listen or not):

The discipline of technology means that the organization must adapt to the technology, and not the technology to the organization. For advanced development, data are never complete, particularly if the product of a complex technology is to operate at high standards for years. The discipline of technology requires exhaustive testing of materials and components to *determine the laws of nature* [emphasis added]. If these are not absolute in the sequestered atmosphere of scientific laboratories or research centers, there is no reason to expect they are better known on the shop floor. The discipline of technology requires thorough and deep consideration of the match between the product and its use, and intense analysis of the present and anticipated future conditions of operation.

This discipline of testing and learning sooner, faster, and cheaper was carried out in many ways. Here is one example: For every version of a shipboard reactor, there was a land-based version on which people could train and on which design problems could be worked out in a safer, cheaper,



more controlled environment. For each such land-based model, there were full-scale wooden and cardboard mock-ups to preview how people and machines would interact in practice. Another example was the testing to see how equipment would handle the shocks of military use. Scaled-down components were mounted in scaled-down submarine hulls with an array of sophisticated gadgetry. Where and how to affix radiation shielding was always a challenging problem. For instance, it might have been easy to calculate the exposure on one side of a smooth rounded surface, but what about convoluted surfaces? When calculation failed, experimentation was the answer; for instance, building a prototype inside a water tank to see what would happen when radiation was emitted out the bottom of the boat but was reflected back into the vessel by the water. (In later chapters we'll see how Toyota makes just such a commitment to use high-speed, low-cost pilots and trials.)

Problem solving within NR has not only been disciplined in terms of the detail, but disciplined in terms of inviting all relevant data and multiple perspectives to a problem. Duncan notes that when shipyard representatives would raise problems with NR headquarters, they would sometimes illustrate the problem with a diagram or a mock-up of a component. A few of [Rickover's] engineers would take over the conference room just outside his office, and when all was ready, he would come in. At the slightest indication of vagueness or ambiguity, he would interrupt, demanding clarity and facts.

The point was not that Rickover always knew better—quite the opposite. The purpose of having detailed write-ups, diagrams, and models was to ensure that competing and



complementary views were well represented. Rank, personality, and assertiveness were not going to determine a decision. The data, coupled with the best collective understanding of a situation, would do that. The NASA benchmarking study stated this as follows:

Recommendations are prepared independently by the prime contractors and undergo extensive internal reviews by experts in all related technical disciplines. The management and personnel at the two NR laboratories are required to provide their technical recommendations independently without soliciting Headquarters' advance agreement. This ensures that each laboratory retains its responsibility for providing its own technical assessment. Any dissenting/alternate opinions are required to be documented in the recommendation with a discussion of the logic for not implementing them.

The NASA study also noted that NR didn't only want to know what people thought was the right answer. It wanted to be very clear where they were uncertain. Therefore, reports from the laboratories had to include, along with their assessments, a clear discussion of alternate or dissenting assessments. As Rickover had explained years earlier:

One must create the ability in his staff to generate clear, forceful arguments for opposing viewpoints as well as their own. Open discussions and disagreements must be encouraged, so that all sides of an issue will be fully explored. Further, important issues



should be presented in writing. Nothing so sharpens the thought process as writing down one's arguments. Weaknesses overlooked in oral discussion rapidly become painfully obvious on the written page.

### **Capability 3: Spreading Lessons Learned to the Whole Organization**

In high-velocity organizations, people do not learn only for themselves. They learn for their colleagues as well. The experiences of an individual contribute to the expertise of the many. Whatever is learned when a problem is seen, swarmed, and solved right where and when it occurs is incorporated into the scripts and specifications to which it applies. Of course, this can only be done if all the assumptions, expectations, and procedures are explicit and available. It would never work if the new knowledge had to be diffused by word of mouth through a complex workplace, never mind a complex constellation of workplaces that might well be scattered over several continents and oceans. In the U.S. nuclear navy, when a new crew assumes responsibility for a new ship, everything it encounters—the design of the ship, the design of its procedures, the design of problem-identification and problem-solving routines, the training—is derived from the Navy's entire cumulative experience.

John Crawford, who rose to be deputy director of the nuclear propulsion program, and Steven Krahn, who spent 10 years in the program working on maintenance and repair, described the organization-wide benefits of turning local discovery into systemic discovery:



This disciplined, formal engineering approach is pervasive in every phase of activities at Naval Reactors: development of codes and standards where none exist, the availability of formalized design manuals and engineered test procedures. . . . [O]ver the years, a comprehensive set of standards and procedures has been developed that has contributed importantly to the safety and reliability of the reactor plants that Naval Reactors builds. This set of standards and procedures permits innovation to be applied in a controlled manner and allows focus to be placed on truly important areas, while ensuring that routine work gets done competently.

The NASA benchmarking team likewise noted that when reports are completed, they are “also provided to other organizations in the program so that they may also learn and take preventive action. This tool has contributed to a program philosophy that underscores the smaller problems in an effort to prevent significant ones.”


Figure 5-3, taken from the NASA benchmarking study, diagrams this constant building of knowledge out of experience, leading to better experiences going forward.

## **Capability 4: Leading by Developing Capabilities 1, 2, and 3 in Others**

What is a leader’s job? It’s common to say that the leader sets goals by dint of his or her greater authority and wider per-



**Figure 5-3 Technical requirements, implementation experience, and lessons learned closed loop**



spective, decides how scarce resources will be allocated among competing priorities, and sets the emotional tone of a particular situation. Certainly Rickover was in the position and of the temperament to do all these.

But in high-performing organizations, the leader has two other critical roles. He or she is responsible for determining not only *what* gets done, by setting goals and allocating resources, but also *how* things get done, by shaping the company's processes and systems. Of course this isn't a one-man or one-woman job, and no one can be leader forever, so he or she must also develop in others the skills needed to lead complex operations.

Rickover modeled both roles. For one thing, he conspicuously modeled the role of leader as learner-in-chief. Ted Rockwell recalls meeting Rickover, soon after World War II, in Oak Ridge, Tennessee, where Rickover had come to learn about nuclear technology on behalf of the Navy. At that time,



Rickover was a captain, a senior officer—the equivalent of a full colonel in the other military services—but Rockwell recalls him as:

. . . this one silver-haired guy who kept asking simple, basic questions, making himself look pretty stupid and getting a lot of knowing chuckles from the wiseacres. That course was pretty tough, even for me, and when one student asked timidly, “Please, Professor, could you tell us, at what level will this course be given?” the prof answered genially, “Let us say at a popular, postdoctoral level.”

At this point, the silver-haired Captain said, “I’m not getting this. Would you please go over it again?” . . . The rest of the class was getting a little restless and wondering why the Navy would send somebody down who was incapable of getting the material. The prof then asked condescendingly, “Would you perhaps like to have us provide you with some tutoring in the evenings?” Not taking this as a put-down, the Captain said merely, “I would appreciate that very much, sir.”

When the remedial session was arranged, it was attended not just by Rickover and the few other Navy personnel but also by the other students, including those who had mocked Rickover’s questions. Upon arriving, he commented: “I guess I’m not the only dummy in the class.”

Rickover was committed not only to his own learning but to others’ learning. Early in his career, while serving as an engi-





neering officer on the battleship *New Mexico*, Rickover demanded that his officers show detailed knowledge of the equipment and machinery for which they were responsible, “calmly accepting mistakes and errors when honestly acknowledged, and giving each man as much responsibility as he could handle.” Three of his ensigns (the lowest rank for a naval officer) went on to become admirals.

Rickover’s commitment to developing his people became institutionalized in the NR program. He personally interviewed every officer candidate for an engineering, construction, or maintenance role or for a position onboard a ship. As notorious as those interviews were for the stress and strain they produced, the objective was not sadism, hazing, or harassment. Rickover wanted to see how people handled pressure, responded to unfamiliar situations, and thought through problems—all of which would determine how well the candidate could manage NR’s complex operations for very high performance. He didn’t necessarily expect a “right” answer, but there were definitely “wrong” answers. “No excuse, sir” for bad grades was worse than a forthright “I was lazy.” A candidate was better off acknowledging that he had reached an illogical or false position than trying to demonstrate conviction by “sticking to his guns.” One candidate’s interview consisted largely of the accusation that he was fat, with the challenge of what he was going to do about it. He responded by detailing changes in his food intake and activity level, with predications of how the change in calories in/calories out would affect his weight. Over the next several weeks, Rickover followed up with the candidate, who both slimmed down and was accepted into the program.



Selection into the NR program triggered an extended, intense learning process. There was the initial six months of on-the-job training in Washington, with an assignment to the office of a more experienced project officer because “new engineers lacked the background to contribute anything.” But it was not enough for the new engineer to learn a particular perspective, role, and skill set in the project office; he was also being introduced to the entire system of nuclear propulsion as well, so he could understand how the piece with which he was becoming familiar fit into the whole. This was accomplished with instruction in “nuclear fission and reactor physics, reactor-plant operations, reactor-core materials, reactor-core design and construction, electrical power systems and instrumentation systems, primary and secondary fluid systems, water chemistry control, radiological control and reactor protection and safety.”

Additional training followed, with the curriculum divided into intense compressed increments. Again, there was the combination of hands-on experience—operating the land-based versions of shipboard reactors—and additional classes in reactor theory and design that laid out not only how these devices were built and operated, but also the reasons why. After running a prototype reactor, officers and crew requalified onboard their ship before beginning two to three years of sea duty. This approach was repeated again for additional progressions in rank and responsibility.

Ultimately, Rickover was driving toward developing a cadre who understood “the discipline of engineering,” the approaches necessary for managing the design, operation, and improvement of systems of great complexity, of great benefit when run well, and of great consequence were they to fail.



Rockwell reflected on his own relationship with Rickover, which extended beyond Rockwell's service in the nuclear-reactor program and many years into his private practice as an engineer. He wrote:

To be categorized as “close” to Rickover needs explanation. It means that he would continue to treat me as his pupil, one still worthy of placing demands on his time and energy to help improve me professionally. This never-ending process of educating and training prospective leaders for the Navy was a driving passion of Rickover's life.

Rickover's leadership, imperious as it may sometimes have been, was a constant refutation of the view that leadership means “to command—someone else would take care of the ship.” By embodying Capability 4, developing highly disciplined problem-identification and problem-solving skills throughout his organization, he ensured that NR would remain a high-velocity organization even without him.

---

### LOSS OF THE *THRESHER*

The nuclear navy's record in submarine safety is not perfect. For example, on April 10, 1963, the *USS Thresher* was lost 200 miles off the coast of the United States, killing the 129 people onboard. Although it may never be possible to know exactly what happened to the *Thresher*, underwater communication, other sensing data, and examination of the wreckage led to the conclusion that it



was felled by flooding. Although the technical details differ from the NASA cases, the same organizational faults that plagued NASA—tolerance for ambiguity in expectations and procedures coupled with willingness to work around obvious problems and to normalize deviance—plagued the branch of the Navy responsible for designing, constructing, and maintaining the nonnuclear portions of its nuclear submarines. In other words, the loss of the *USS Thresher* for non-reactor reasons makes for a striking contrast between high-velocity and low-velocity organizations and illustrates how both approaches can exist in the same parent organization if great care is not exercised by leadership.

Let's take a closer look at what happened. On a submarine, a leak in a pipe can be catastrophic. Because of the intense pressure in the lines when the ship is submerged, even a small leak that would be an annoying drip-drip on the surface can create a blinding spray, incapacitating the crew and shorting out electrical equipment. A large leak can flood a vessel, making resurfacing impossible. Therefore, when it comes to running pipes through submarines, the quality of the welds is paramount. In the late 1950s and early 1960s, there were two ways of joining pipes: welding and silver brazing. Although silver brazing was perceived to have advantages when done properly, it was technically more difficult and experience suggested that it was not reliable enough.

For example, in November 1960, the submarine *USS Barb* left Norfolk to participate in an exercise with other



ships. Its captain began a series of test dives, leveling off every 100 feet to check that nothing was amiss. At test depth, there was a report of flooding in the engine room. The crew sealed the flooded compartment, the ballast was blown, and the engines were set to full speed ahead as the captain successfully drove the ship to the surface. A pipe carrying salt water had given way at a silver-brazed joint. Other silver-brazed joints later failed on the submarine *USS Abraham Lincoln*. Inspections on yet another submarine revealed poorly brazed joints. Nevertheless, the Navy proceeded to build and operate submarines with those joints in critical lines. On another test dive a small saltwater line failed, and other pipe failures were documented.

As for the *Thresher*, it had 3,000 silver-brazed joints that were subject to full pressure. During a maintenance inspection, 145 of them were inspected, with 14 percent showing irregularities. That rate across all the joints would have meant over 400 joints with possible defects, yet the ship was put out to sea. It shouldn't have been.

As an accompanying surface ship listened through underwater devices, the *Thresher's* crew encountered some difficulty, tried to surface, couldn't, and sank. Based on the sound recordings that were made at the time, the accident investigation report surmised that a pipe had burst, the crew had been unable to stop the flooding, and spray had short-circuited equipment, causing a loss of power. The ballast tanks did not operate properly, so the submarine continued to sink and finally was crushed by the pressure.



Prompted by that catastrophe, the Navy created its SUBSAFE program, an approach to designing, constructing, and maintaining the nonnuclear portions of the submarine (responsible for submergence, flood prevention, control, and recovery) as disciplined and rigorous as the NR approach.

---

Let's look at two other organizations that made the transition to high velocity. One, Avenue A (later known as Quantive), was a pioneer in managing online advertising campaigns. The other, Pratt & Whitney, designs and builds commercial and military jet engines. These two organizations could hardly differ more in their markets, customers, suppliers, and technology, but both offer complex products developed in competitive industries that depend on the most advanced science and engineering. And both dramatically increased their performance by deciding to manage the functional pieces of their enterprises holistically as parts of a well-integrated whole and by recognizing that achieving high performance depended on building deeper system knowledge and could never be accomplished through workarounds and firefighting.

### **Pratt & Whitney: Higher-Speed, Lower-Cost New Product Development**


The jet engine is a technological marvel, a vast improvement over the piston-driven propeller engines that dominated avia-



tion through the end of the 1940s. Able to generate far more power than previously possible, jet engines have revolutionized air travel and have found use as power plants on land as well. The mechanics of a jet engine are conceptually simple—a basic application of the principle that for every action, there is an equal and opposite reaction. Throw a medicine ball in one direction while wearing roller skates and you are bound to roll away in the other direction.


In the case of the jet engine, this is what happens. A mixture of fuel and air are ignited in a combustion chamber to generate thrust, as shown in Figure 5-4. Part of the thrust drives the engine (and the plane to which it is attached) forward. Part of the thrust is captured by the turbine fan blades behind the engine. These turbines are connected by a drive shaft and spin the compressor blades at the front of the engine, which drive more air, under high pressure, into the combustion chamber to feed the process. The turbofan variation on this approach is to have two sets of turbines on the back of the engine, as shown in Figure 5-5. One drives the high-pressure compressor blades, pumping air into the com-

**Figure 5-4 Jet engine basics**



Source: NASA Web site, “Types of Gas Turbines, Glenn Research Center,” at <http://www.grc.nasa.gov/WWW/K-12/airplane/trbtyp.html>.



**Figure 5-5 Turbofan jet engine**

*Source:* NASA Web site, “Types of Gas Turbines, Glenn Research Center,” at <http://www.grc.nasa.gov/WWW/K-12/airplane/trbtyp.html>.

bustion chamber. The other is connected to low-pressure compressor blades which move massive volumes of air past the combustion chamber and directly out the back of the engine to produce thrust. The advantages of this approach include fuel efficiency and quieter operation.

All this required extraordinary advances in materials science, combustion dynamics, aerodynamics, and control systems. Each of these is an extremely complex discipline in its own right; getting them to work together is no trifle. As a result, jet-engine compressors, combustion chambers, turbines, and other components—many of which have to work extremely reliably in conditions of extreme heat, pressure, and stress—are the products of intense research and development efforts. To ensure that they all work in concert, there have been advances in aerodynamics and fluid dynamics within the various parts of the engine and advances in the electronic controls, gearing, bearings, and virtually every other part of the engine.

In its early years, Pratt had enough time and money to figure all this out as it went along, taking a “think tank”





approach to technology development. Defense contracts were rich sources of revenue and commercial contracts, once won, were a predictable income stream. After an engine manufacturer won the rights to supply the engine for a particular model of commercial aircraft, every plane of that type would be delivered with that engine. This, in turn, guaranteed the engine manufacturer a steady annuity for spare parts and other maintenance services. In those days, a young engineer joined what was essentially an apprenticeship. One engineer explained that when he first started at Pratt, his boss would give him instructions, which he would carry out. Then his boss would check his work, have him fix the problems, and check his work again. When it was deemed acceptable, it went to the next boss for checking. Quality came from hard work, inspection, and rework. If a problem couldn't be resolved on paper, engineers could build a prototype, run it until it broke down—a method that came to be known as “build and bust”—and put whatever knowledge was discovered to use in the next iteration.

But as times changed, this approach became untenable. Not only did military spending decrease with the end of the Cold War, but the commercial market changed as well. Airline deregulation increased the price-sensitivity of airlines and new entrants to the market increased the competition. Airplane makers began to certify more than one engine per airliner model. An engine maker now had to woo not only Boeing and Airbus, but each of the airlines as well. Of course, more competition and less customer lock-up made for unprecedented cost pressures on the design and development, manufacturing, and spare-parts support of jet engines.



Pratt could not afford its think-tank approach any longer. Developing a new engine had been up to a \$1 billion investment, requiring nearly four years of work, but the new market dynamics convinced management that development costs had to be cut to \$300 million per platform and that development time had to be cut down to 30 months.

To achieve these goals, Pratt created integrated program-management teams in the early 1990s. The rationale was that creating cross-functional teams, with representatives from different disciplines and different components working together, would reduce the rework and expense that came from developing complex jet-engine components in isolation. That was an improvement, inasmuch as it addressed a key structural shortcoming of low-velocity organizations—managing the pieces of their systems without an eye to how the pieces need to come together.

But it was not good enough. Pratt was still missing the dynamic component of high-velocity management—generating useful knowledge and building on it rather than having to keep acquiring the same knowledge over and over. A 50 percent downsizing only made matters worse as Pratt lost some of its most experienced engineers and managers, who not only had deep technical knowledge in their own realms but also had acquired knowledge about ways to navigate and coordinate the work system to achieve good outcomes.

Pratt now needed a better way to ensure that (a) when people started their work, they could bring to bear the cumulative experience of the whole organization, and (b) when the cumulative expertise of the organization was found to be missing something, insights from new experiences



could quickly become part of that available cumulative experience. Pratt's attempt to do this was called "engineering standard work" (ESW), making a true discipline of its engineering efforts much as the Navy's nuclear reactor program had done.

First, Pratt engineers laid out everything that was already known about the design process. Extensive *workflow maps* made clear what design steps normally occurred in what sequence and with what interdependence. Understanding what step was dependent on what other steps set the stage for establishing *design criteria* to clarify what each step had to accomplish to satisfy the needs of those who depended on it. To increase the chance that each stage would be successful in meeting those criteria, *activity pages* were created, representing the best method for achieving success known at the time, with tools and methods instructions explaining how, when, and why various analytical and other design tools should be used. Then, to determine who was capable of being responsible for what stage, with what degree of support and supervision, Pratt created *practitioner proficiency assessments* to determine how much support someone needed in a role or how much he or she could provide. Readiness reviews determined if a new technology could be mainstreamed into a program or if it was still developmental.

These were all mechanisms for capturing and sharing knowledge. There was also a mechanism for building knowledge. When someone encountered a problem while using some element of ESW—a workflow map, design criteria, or an activity page—there was an owner of that element who could be called in to investigate. When the root cause was dis-



covered, the ESW was modified, increasing the likelihood that the next person to depend on that element of ESW would succeed.

Paul Adams, a longtime engineer at Pratt, explained to me the effect of making sure that local learning became organizational learning:

First, we had to make sure we had the handoffs down, controlling how you work in a dispersed organization. The workflow maps, design criteria, and all of that picked up that piece. We also had to get a handle on how to use the new computational tools. They are very useful, but only within boundaries that have been proven. Outside those boundaries, you're taking some real risks. We had to give people clarity as to the situations in which those tools worked and the situations in which they didn't.

(In contrast, NASA hadn't done such a good job of that, leaving it to an inexperienced engineer to use an unfamiliar software package outside its design specs to estimate the impact damage on the wing of the *Columbia*.)

Another concern that led us to standard work was: How quickly could somebody be effective? Work flow maps, design criteria, proficiency tests—those are about execution of standard work. If we can teach and test skills, [it creates] a very substantial decrease in the time it would take people to get proficient.



The first thing people say, before they are really exposed to ESW, is that this takes away my ability to innovate. But that is wrong. It gives you a chance to innovate in a controlled manner, so you won't introduce additional risk into the product. With standards, we can distinguish where we don't need you to innovate from where you need to innovate. It helps you see where innovation is needed and helps you determine what innovation is useful. When we do need to innovate more quickly, standard work helps with our technology readiness process, because better clarity about what we have and how it fits into the entirety of the program helps us understand where we have a high and low risk tolerance.

The results of always using the best approach that had been found so far, of making sure that the people given responsibility for a task were actually capable of it, and of making sure that local improvements became organizational improvements were quite good. The commercial and military projects on which ESW was piloted came in on time and on budget—hardly a familiar experience. Engineering change orders, those late-stage design changes that are costly to implement because so much is already set in place, were down by half in the first year of using ESW and down another 15 percent in the second year. That alone saved an estimated \$50 million on rework. All told, Pratt estimated that every dollar spent on ESW yielded a four-dollar payback.



## Avenue A: From Chaotic Mess to Lean and Mean

Avenue A was a pioneer in creating Web-based marketing. The advantages it offered its clients were immense. Advertisers could target their audiences with far greater precision than they could with television, radio, telephone, or direct mail. They could get immediate feedback, finding out who was responding in what fashion to what ad on what Web page with what frequency, and modify their advertising quickly for maximum impact. In 1999, after only three years of operation, business was booming, with revenue growth of 50 percent *per month*, but the firm's poor work processes were getting in the way of its success. Head count kept increasing, but with no appreciable increase in efficiency. Employees were putting more and more effort into scheduling, coordinating, clarifying, and redoing their work and less into designing, implementing, and optimizing marketing campaigns. More success actually meant bigger losses.

Eight years and one disastrous dot-com bubble burst later, Avenue A not only had survived, but was flourishing. Now known as aQuantive, it had grown from three employees to over 2,000. In mid-2007, Microsoft bought aQuantive for \$6 billion, quite a return on the early investment of \$20.5 million.

How did Avenue A pull this off? There was no way for it to control an external environment as fluid as the Web, so it had to shape the fluidity of its internal environment with great sophistication, improving and innovating more quickly, for longer durations, and with greater breadth, in order to set itself apart. Let's take a close look at how this was accomplished.

Avenue A grew out of Nick Hanauer's efforts to advertise his family's business, Pacific Coast Feather Company. To reach a



wider customer base, he built an online catalog, but soon found that if people didn't know about the site, they didn't visit it. He then had the idea of buying advertising space on other Web sites, paying a commission for every click-through. As his experience grew, Hanauer had the idea of brokering extra space that he didn't need. Then, as his expertise grew in designing marketing campaigns that could be managed in real time for Pacific Coast Feather, he realized that there was a business in helping others do the same. Avenue A was created as a stand-alone business offering three services:

1. *Design.* Planning an Internet-based media campaign and buying ad space (which included negotiating the rates).
2. *Implementation.* Providing the technical support for an advertising campaign: housing the ads on Avenue A's own servers and placing them according to the campaign's plan.
3. *Optimization.* Gathering data about which ads on which sites led to reader click-throughs and using those data to modify the advertising campaign.

As straightforward as that seems, there were many steps within each stage and, as Avenue A grew, an increasing number of people responsible for each step. For instance, within the design phase, an Avenue A representative had to work with the client to develop the themes and approaches of the campaign. Someone else, with expertise in various types of publishers, had to identify what type of Web sites would be most promising for a particular type of ad. These suggestions would have to be run past the client for approval before Avenue A could go back to the publishers to negotiate rates. This never



got easier as the number of publishers increased and advertising options multiplied. What began as a fairly simple system—a few people doing a few things—grew in complexity, with many activities dependent on each other in often surprising ways. The results were predictable—delays, defects that had to be caught on the fly, missed handoffs, and a general demand that people go above and beyond, all the time, to succeed.

Avenue A's initial efforts to offset these problems were understandable, if not effective. Where there were bottlenecks, work was shifted from one group of specialists to another. If that didn't do it, more people were added. But that never solved the basic problem: The pieces didn't come together well without heroic efforts on everyone's part. Projects ricocheted around the organization, repeatedly ping-ponging between someone who needed something and the person who had worked hard, but not successfully, to provide it. For example, someone might have designed a marketing campaign—what ads would go on what publisher pages. Those specs would be passed to someone else who would then discover that some of those publishers couldn't support the types of ads that had been specified. Once that problem had been resolved—after several rounds of ping-ponging e-mails and phone calls—instructions would go to the engineer responsible for implementing the campaign on Avenue A's servers. He would discover that the computer codes and protocols he had been given didn't work. Once again, the work ricocheted from one person to another for clarification, modification, and renovation.

When the unmanageability reached a tipping point, Avenue A divided itself into eight teams, each a stand-alone micro-





cosm of the larger organization with the full complement of specialties needed to pull off a campaign. That didn't eliminate the problems, it just cloned them.

Finally, Avenue A stepped back from the madness and began to work out a system for building knowledge as it solved problems rather than working around the same problems day in and day out. First, it mapped out all the work that had to be done to move a campaign from concept to completion. Then, it determined what each step had to accomplish for its work to be usable by the next step. For example, if computer codes were provided, what did it mean for them to be correct from the user's perspective? Avenue A's improvisational ways of converting a concept into reality were standardized and automated. This meant that innovative energy could be directed to devising better approaches, not to coping with flawed ones. For example, Avenue A's collective knowledge about publishers was extensive, but split up among individual media buyers who each had his or her personal expertise. The whole was much less than the sum of the parts because so much effort had to be wasted in finding out who knew the important fact of the moment. An investment was made to collect all that information so that when someone learned something new about a particular publisher, it quickly became part of everyone's expertise.

With this accumulation of expertise came the profitable opportunity to break the business into distinct modules rather than having everything tied to everything else through the tangle of requests and responses. For example, one group was expert at working with clients on campaign design. Another specialized in publisher relations, building a knowledge base



about which sites were appropriate for which target audiences, how the sites had performed on previous campaigns, what was available, and what the rates would be. A third built technical expertise for campaign implementation and optimization. The services of the three distinct modules could be sold individually or in combination. A technically savvy organization could pay for design consultation. A traditional marketing firm could do its own design and receive media-buying services and technical support. All three services were now effective and reliable, both individually and collectively. Avenue A had become a high-velocity organization, managing its complex operations to deliver a complex service with a high and always-increasing level of performance.

## And Now for Toyota . . .

In the next section, we are going to take a closer look at the four capabilities we have been discussing all along. Chapters 6 through 10 are based on my opportunity to observe and experience what form these four capabilities take and how they are managed and perpetuated at one of the all-time high-velocity organizations, Toyota. Chapter 6 will show how Toyota's work systems are designed not only to capture the best currently known approaches to lead to success but also to reveal deficiencies in systems design when and where they occur. Chapter 7 will show how problem solving is practiced by people at various levels of responsibility and how these skills are inculcated. Chapter 8 will show how learning acquired locally by individuals and small groups is converted into collective knowledge for the



entire organization. Chapter 9 will reveal what kind of leadership is required in a high-velocity organization. Chapter 10 will present several cases showing how Toyota engages these four capabilities not just for routine repeatable work but for large-scale, one-off recoveries from crises. Any reader who still thinks the approaches described in this book are for incremental change only will have his or her perception changed by the end of Chapter 10. Then, before concluding, Chapter 11 will look at health-care delivery organizations that have used lessons from Toyota to help more people and harm fewer, all the while working less hard and at reduced cost.



