

Building a Learning Organization

by David A. Garvin

FROM THE JULY–AUGUST 1993 ISSUE

Continuous improvement programs are sprouting up all over as organizations strive to better themselves and gain an edge. The topic list is long and varied, and sometimes it seems as though a program a month is needed just to keep up. Unfortunately, failed programs far outnumber successes, and improvement rates remain distressingly low. Why? Because most companies have failed to grasp a basic truth. Continuous improvement requires a commitment to learning.

How, after all, can an organization improve without first learning something new? Solving a problem, introducing a product, and reengineering a process all require seeing the world in a new light and acting accordingly. In the absence of learning, companies—and individuals—simply repeat old practices. Change remains cosmetic, and improvements are either fortuitous or short-lived.

A few farsighted executives—Ray Stata of Analog Devices, Gordon Forward of Chaparral Steel, Paul Allaire of Xerox—have recognized the link between learning and continuous improvement and have begun to refocus their companies around it. Scholars too have jumped on the bandwagon, beating the drum for “learning organizations” and “knowledge-creating companies.” In rapidly changing businesses like semiconductors and consumer electronics, these ideas are fast taking hold. Yet despite the encouraging signs, the topic in large part remains murky, confused, and difficult to penetrate.

Meaning, Management, and Measurement

Scholars are partly to blame. Their discussions of learning organizations have often been reverential and utopian, filled with near mystical terminology. Paradise, they would have you believe, is just around the corner. Peter Senge, who popularized learning organizations in his book *The Fifth Discipline*, described them as places “where people continually expand their capacity to create the results they truly desire, where new and expansive patterns of thinking are nurtured, where collective aspiration is set free, and where people are continually learning how to learn together.”¹ To achieve these ends, Senge suggested the use of five “component technologies”: systems thinking, personal mastery, mental models, shared vision, and team learning. In a similar spirit, Ikujiro Nonaka characterized knowledge-creating companies as places where “inventing new knowledge is not a specialized activity...it is a way of behaving, indeed, a way of being, in which everyone is a knowledge worker.”² Nonaka suggested that companies use metaphors and organizational redundancy to focus thinking, encourage dialogue, and make tacit, instinctively understood ideas explicit.

Sound idyllic? Absolutely. Desirable? Without question. But does it provide a framework for action? Hardly. The recommendations are far too abstract, and too many questions remain unanswered. How, for example, will managers know when their companies have become learning organizations? What concrete changes in behavior are required? What policies and programs must be in place? How do you get from here to there?

Most discussions of learning organizations finesse these issues. Their focus is high philosophy and grand themes, sweeping metaphors rather than the gritty details of practice. Three critical issues are left unresolved; yet each is essential for effective implementation. First is the question of *meaning*. We need a plausible, well-grounded definition of learning organizations; it must be actionable and easy to apply. Second is the question of *management*. We need clearer guidelines for practice, filled with operational advice rather than high aspirations. And third is the question of *measurement*. We need better tools for assessing an organization's rate and level of learning to ensure that gains have in fact been made.

Once these “three Ms” are addressed, managers will have a firmer foundation for launching learning organizations. Without this groundwork, progress is unlikely, and for the simplest of reasons. For learning to become a meaningful corporate goal, it must first be understood.

What Is a Learning Organization?

Surprisingly, a clear definition of learning has proved to be elusive over the years. Organizational theorists have studied learning for a long time; the accompanying quotations suggest that there is still considerable disagreement (see the insert “Definitions of Organizational Learning”). Most scholars view organizational learning as a process that unfolds over time and link it with knowledge acquisition and improved performance. But they differ on other important matters.

Definitions of Organizational Learning

Scholars have proposed a variety of definitions of organizational learning. Here is a small sample:

“Organizational learning means the process of improving actions through better knowledge and understanding.” —C. Marlene Fiol and Marjorie A. Lyles, “Organizational Learning,” *Academy of Management Review*, October 1985.

“An entity learns if, through its processing of information, the range of its potential behaviors is changed.” —George P. Huber, “Organizational Learning: The Contributing Processes and the Literatures,” *Organization Science*, February 1991.

“Organizations are seen as learning by encoding inferences from history into routines that guide behavior.” —Barbara Levitt and James G. March, “Organizational Learning,” *American Review of Sociology*, Vol. 14, 1988.

“Organizational learning is a process of detecting and correcting error.” —Chris Argyris, “Double Loop Learning in Organizations,” *Harvard Business Review*, September–October 1977.

“Organizational learning occurs through shared insights, knowledge, and mental models...[and] builds on past knowledge and experience—that is, on memory.” —Ray Stata, “Organizational Learning—The Key to Management Innovation,” *Sloan Management Review*, Spring 1989.

Some, for example, believe that behavioral change is required for learning; others insist that new ways of thinking are enough. Some cite information processing as the mechanism through which learning takes place; others propose shared insights, organizational routines, even memory. And some think that organizational learning is common, while others believe that flawed, self-serving interpretations are the norm.

How can we discern among this cacophony of voices yet build on earlier insights? As a first step, consider the following definition:

A learning organization is an organization skilled at creating, acquiring, and transferring knowledge, and at modifying its behavior to reflect new knowledge and insights.

This definition begins with a simple truth: new ideas are essential if learning is to take place. Sometimes they are created *de novo*, through flashes of insight or creativity; at other times they arrive from outside the organization or are communicated by knowledgeable insiders. Whatever their source, these ideas are the trigger for organizational improvement. But they cannot by themselves create a learning organization. *Without accompanying changes in the way that work gets done, only the potential for improvement exists.*

This is a surprisingly stringent test for it rules out a number of obvious candidates for learning organizations. Many universities fail to qualify, as do many consulting firms. Even General Motors, despite its recent efforts to improve performance, is found wanting. All of these organizations have been effective at creating or acquiring new knowledge but notably less successful in applying that knowledge to their own activities. Total quality management, for example, is now taught at many business schools, yet the number using it to guide their own decision making is very small. Organizational consultants advise clients on social dynamics and small-group behavior but are notorious for their own infighting and factionalism. And GM, with a few exceptions (like Saturn and NUMMI), has had little success in revamping its manufacturing practices, even though its managers are experts on lean manufacturing, JIT production, and the requirements for improved quality of work life.

Organizations that do pass the definitional test—Honda, Corning, and General Electric come quickly to mind—have, by contrast, become adept at translating new knowledge into new ways of behaving. These companies actively manage the learning process to ensure that it occurs by design rather than by chance. Distinctive policies and practices are responsible for their success; they form the building blocks of learning organizations.

Building Blocks

Learning organizations are skilled at five main activities: systematic problem solving, experimentation with new approaches, learning from their own experience and past history, learning from the experiences and best practices of others, and transferring knowledge quickly and efficiently throughout the organization. Each is accompanied by a distinctive mind-set, tool kit, and pattern of behavior. Many companies practice these activities to some degree. But few are consistently successful because they rely largely on happenstance and isolated examples. By creating systems and processes that support these activities and integrate them into the fabric of daily operations, companies can manage their learning more effectively.

1. Systematic problem solving

This first activity rests heavily on the philosophy and methods of the quality movement. Its underlying ideas, now widely accepted, include:

- Relying on the scientific method, rather than guesswork, for diagnosing problems (what Deming calls the “Plan, Do, Check, Act” cycle, and others refer to as “hypothesis-generating, hypothesis-testing” techniques).
- Insisting on data, rather than assumptions, as background for decision making (what quality practitioners call “fact-based management”).
- Using simple statistical tools (histograms, Pareto charts, correlations, cause-and-effect diagrams) to organize data and draw inferences.

Most training programs focus primarily on problem-solving techniques, using exercises and practical examples. These tools are relatively straightforward and easily communicated; the necessary mind-set, however, is more difficult to establish. Accuracy and precision are essential for learning. Employees must therefore become more disciplined in their thinking and more attentive to details. They must continually ask, “How do we know that’s true?”, recognizing that close enough is not good enough if real learning is to take place. They must push beyond obvious symptoms to assess underlying causes, often collecting evidence when conventional wisdom says it is unnecessary. Otherwise, the organization will remain a prisoner of “gut facts” and sloppy reasoning, and learning will be stifled.

Xerox has mastered this approach on a company-wide scale. In 1983, senior managers launched the company’s Leadership Through Quality initiative; since then, all employees have been trained in small-group activities and problem-solving techniques. Today a six-step process is used for virtually all decisions (see the insert “Xerox’s Problem-Solving Process”). Employees are provided with tools in four areas: generating ideas and collecting information

(brainstorming, interviewing, surveying); reaching consensus (list reduction, rating forms, weighted voting); analyzing and displaying data (cause-and-effect diagrams, force-field analysis); and planning actions (flow charts, Gantt charts). They then practice these tools during training sessions that last several days. Training is presented in “family groups,” members of the same department or business-unit team, and the tools are applied to real problems facing the group. The result of this process has been a common vocabulary and a consistent, companywide approach to problem solving. Once employees have been trained, they are expected to use the techniques at all meetings, and no topic is off-limits. When a high-level group was formed to review Xerox’s organizational structure and suggest alternatives, it employed the very same process and tools.³

Xerox’s Problem-Solving Process				
Step	Question to Be Answered	Expansion/Divergence	Contraction/Convergence	What’s Needed to Go to the Next Step
1. Identify and select problem	What do we want to change?	Lots of problems for consideration	One problem statement, one “desired state” agreed upon	Identification of the gap “Desired state” described in observable terms
2. Analyze problem	What’s preventing us from reaching the “desired state”?	Lots of potential causes identified	Key cause(s) identified and verified	Key cause(s) documented and ranked
3. Generate potential solutions	How could we make the change?	Lots of ideas on how to solve the problem	Potential solutions clarified	Solution list
4. Select and plan the solution	What’s the best way to do it?	Lots of criteria for evaluating potential solutions Lots of ideas on how to implement and evaluate the selected solution	Criteria to use for evaluating solution agreed upon Implementation and evaluation plans agreed upon	Plan for making and monitoring the change Measurement criteria to evaluate solution effectiveness
5. Implement the solution	Are we following the plan?		Implementation of agreed-on contingency plans (if necessary)	Solution in place
6. Evaluate the solution	How well did it work?		Effectiveness of solution agreed upon Continuing problems (if any) identified	Verification that the problem is solved, or Agreement to address continuing problems

Xerox’s Problem-Solving Process

2. Experimentation

This activity involves the systematic searching for and testing of new knowledge. Using the scientific method is essential, and there are obvious parallels to systematic problem solving. But unlike problem solving, experimentation is usually motivated by opportunity and expanding horizons, not by current difficulties. It takes two main forms: ongoing programs and one-of-a-kind demonstration projects.

Ongoing programs normally involve a continuing series of small experiments, designed to produce incremental gains in knowledge. They are the mainstay of most continuous improvement programs and are especially common on the shop floor. Corning, for example, experiments continually with diverse raw materials and new formulations to increase yields and provide better grades of glass. Allegheny Ludlum, a specialty steelmaker, regularly examines new rolling methods and improved technologies to raise productivity and reduce costs.

Opportunity motivates experimentation. Corning, for example, continually strives to increase yields and provide better grades of glass.

Successful ongoing programs share several characteristics. First, they work hard to ensure a steady flow of new ideas, even if they must be imported from outside the organization. Chaparral Steel sends its first-line supervisors on sabbaticals around the globe, where they visit academic and industry leaders, develop an understanding of new work practices and technologies, then bring what they’ve learned back to the company and apply it to daily operations. In large part as a result of these initiatives, Chaparral is one of the five lowest cost steel plants in the world. GE’s Impact Program originally sent manufacturing managers to Japan to study factory innovations, such as quality circles and

kanban cards, and then apply them in their own organizations; today Europe is the destination, and productivity improvement practices the target. The program is one reason GE has recorded productivity gains averaging nearly 5% over the last four years.

Successful ongoing programs also require an incentive system that favors risk taking. Employees must feel that the benefits of experimentation exceed the costs; otherwise, they will not participate. This creates a difficult challenge for managers, who are trapped between two perilous extremes. They must maintain accountability and control over experiments without stifling creativity by unduly penalizing employees for failures. Allegheny Ludlum has perfected this juggling act: it keeps expensive, high-impact experiments off the scorecard used to evaluate managers but requires prior approvals from four senior vice presidents. The result has been a history of productivity improvements annually averaging 7% to 8%.

Successful programs require an incentive system that favors risk taking.

Finally, ongoing programs need managers and employees who are trained in the skills required to perform and evaluate experiments. These skills are seldom intuitive and must usually be learned. They cover a broad sweep: statistical methods, like design of experiments, that efficiently compare a large number of alternatives; graphical techniques, like process analysis, that are essential for redesigning work flows; and creativity techniques, like storyboarding and role playing, that keep novel ideas flowing. The most effective training programs are tightly focused and feature a small set of techniques tailored to employees' needs. Training in design of experiments, for example, is useful for manufacturing engineers, while creativity techniques are well suited to development groups.

Demonstration projects are usually larger and more complex than ongoing experiments. They involve holistic, systemwide changes, introduced at a single site, and are often undertaken with the goal of developing new organizational capabilities. Because these projects represent a sharp break from the past, they are usually designed from scratch, using a "clean slate" approach. General Foods's Topeka plant, one of the first high-commitment work systems in this country, was a pioneering demonstration project initiated to introduce the idea of self-managing teams and high levels of worker autonomy; a more recent example, designed to rethink small-car development, manufacturing, and sales, is GM's Saturn Division.

Demonstration projects share a number of distinctive characteristics:

- They are usually the first projects to embody principles and approaches that the organization hopes to adopt later on a larger scale. For this reason, they are more transitional efforts than endpoints and involve considerable "learning by doing." Mid-course corrections are common.
- They implicitly establish policy guidelines and decision rules for later projects. Managers must therefore be sensitive to the precedents they are setting and must send strong signals if they expect to establish new norms.
- They often encounter severe tests of commitment from employees who wish to see whether the rules have, in fact, changed.
- They are normally developed by strong multi-functional teams reporting directly to senior management. (For projects targeting employee involvement or quality of work life, teams should be multilevel as well.)
- They tend to have only limited impact on the rest of the organization if they are not accompanied by explicit strategies for transferring learning.

All of these characteristics appeared in a demonstration project launched by Copeland Corporation, a highly successful compressor manufacturer, in the mid-1970s. Matt Diggs, then the new CEO, wanted to transform the company's approach to manufacturing. Previously, Copeland had machined and assembled all products in a single facility. Costs were high, and quality was marginal. The problem, Diggs felt, was too much complexity.

At the outset, Diggs assigned a small, multifunctional team the task of designing a "focused factory" dedicated to a narrow, newly developed product line. The team reported directly to Diggs and took three years to complete its work. Initially, the project budget was \$10 million to \$12 million; that figure was repeatedly revised as the team found, through experience and with Diggs's prodding, that it could achieve dramatic improvements. The final investment, a total of \$30 million, yielded unanticipated breakthroughs in reliability testing, automatic tool adjustment, and programmable control. All were achieved through learning by doing.

The team set additional precedents during the plant's start-up and early operations. To dramatize the importance of quality, for example, the quality manager was appointed second-in-command, a significant move upward. The same reporting relationship was used at all subsequent plants. In addition, Diggs urged the plant manager to ramp up slowly to full production and resist all efforts to proliferate products. These instructions were unusual at Copeland, where the marketing department normally ruled. Both directives were quickly tested; management held firm, and the implications were felt throughout the organization. Manufacturing's stature improved, and the company as a whole recognized its competitive contribution. One observer commented, "Marketing had always run the company, so they couldn't believe it. The change was visible at the highest levels, and it went down hard."

Once the first focused factory was running smoothly—it seized 25% of the market in two years and held its edge in reliability for over a decade—Copeland built four more factories in quick succession. Diggs assigned members of the initial project to each factory's design team to ensure that early learnings were not lost; these people later rotated into operating assignments. Today focused factories remain the cornerstone of Copeland's manufacturing strategy and a continuing source of its cost and quality advantages.

Whether they are demonstration projects like Copeland's or ongoing programs like Allegheny Ludlum's, all forms of experimentation seek the same end: moving from superficial knowledge to deep understanding. At its simplest, the distinction is between knowing how things are done and knowing why they occur. Knowing how is partial knowledge; it is rooted in norms of behavior, standards of practice, and settings of equipment. Knowing why is more fundamental: it captures underlying cause-and-effect relationships and accommodates exceptions, adaptations, and unforeseen events. The ability to control temperatures and pressures to align grains of silicon and form silicon steel is an example of knowing how; understanding the chemical and physical process that produces the alignment is knowing why.

Further distinctions are possible, as the insert "Stages of Knowledge" suggests. Operating knowledge can be arrayed in a hierarchy, moving from limited understanding and the ability to make few distinctions to more complete understanding in which all contingencies are anticipated and controlled. In this context, experimentation and problem solving foster learning by pushing organizations up the hierarchy, from lower to higher stages of knowledge.

Stages of Knowledge

Scholars have suggested that production and operating knowledge can be classified systematically by level or stage of understanding. At the lowest levels of manufacturing knowledge, little is known other than the characteristics of a good product. Production remains an art, and there are few clearly articulated standards or rules. An example would be Stradivarius violins. Experts agree that they produce vastly superior sound, but

3. Learning from past experience

Companies must review their successes and failures, assess them systematically, and record the lessons in a form that employees find open and accessible. One expert has called this process the "Santayana Review," citing the famous philosopher George Santayana, who coined the phrase "Those who cannot remember the past are condemned to

no one can specify precisely how they were manufactured because skilled artisans were responsible. By contrast, at the highest levels of manufacturing knowledge, all aspects of production are known and understood. All materials and processing variations are articulated and accounted for, with rules and procedures for every contingency. Here an example would be a “lights out,” fully automated factory that operates for many hours without any human intervention.

In total, this framework specifies eight stages of knowledge. From lowest to highest, they are:

1. Recognizing prototypes (what is a good product?).
2. Recognizing attributes within prototypes (ability to define some conditions under which process gives good output).
3. Discriminating among attributes (which attributes are important? Experts may differ about relevance of patterns; new operators are often trained through apprenticeships).
4. Measuring attributes (some key attributes are measured; measures may be qualitative and relative).
5. Locally controlling attributes (repeatable performance; process designed by expert, but technicians can perform it).
6. Recognizing and discriminating between contingencies (production process can be mechanized and monitored manually).
7. Controlling contingencies (process can be automated).
8. Understanding procedures and controlling contingencies (process is completely understood).

repeat it.” Unfortunately, too many managers today are indifferent, even hostile, to the past, and by failing to reflect on it, they let valuable knowledge escape.

A study of more than 150 new products concluded that “the knowledge gained from failures [is] often instrumental in achieving subsequent successes... In the simplest terms, failure is the ultimate teacher.”⁴ IBM’s 360 computer series, for example, one of the most popular and profitable ever built, was based on the technology of the failed Stretch computer that preceded it. In this case, as in many others, learning occurred by chance rather than by careful planning. A few companies, however, have established processes that require their managers to periodically think about the past and learn from their mistakes.

Boeing did so immediately after its difficulties with the 737 and 747 plane programs. Both planes were introduced with much fanfare and also with serious problems. To ensure that the problems were not repeated, senior managers commissioned a high-level employee group, called Project Homework, to compare the development processes of the 737 and 747 with those of the 707 and 727, two of the company’s most profitable planes. The group was asked to develop a set of “lessons learned” that could be used on future projects. After working for three years, they produced hundreds of recommendations and an inch-thick booklet. Several members of the team were then transferred to the 757 and 767 start-ups, and guided by experience, they produced the most successful, error-free launches in Boeing’s history.

Boeing used lessons from earlier model development to help produce the 757 and 767—the most successful, error-free launches in its history.

Other companies have used a similar retrospective approach. Like Boeing, Xerox studied its product development process, examining three troubled products in an effort to understand why the company’s new business initiatives failed so often. Arthur D. Little, the consulting company, focused on its past successes. Senior management invited ADL consultants from around the world to a two-day “jamboree,” featuring booths and presentations documenting a wide range of the company’s most successful practices, publications, and techniques. British Petroleum went even further and established the post-project appraisal unit to review major investment projects, write up case studies, and derive lessons for planners that were then incorporated into revisions of the company’s planning guidelines. A five-person unit reported to the board of directors and reviewed six projects annually. The bulk of the time was spent in the field interviewing managers.⁵ This type of review is now conducted regularly at the project level.

At the heart of this approach, one expert has observed, “is a mind-set that...enables companies to recognize the value of productive failure as contrasted with unproductive success. A productive failure is one that leads to insight, understanding, and thus an addition to the commonly held wisdom of the organization. An unproductive success occurs when something goes well, but nobody knows how or why.”⁶ IBM’s legendary founder, Thomas Watson, Sr., apparently understood the distinction well. Company lore has it that a young manager, after losing \$10 million in a risky venture, was called into Watson’s office. The young man, thoroughly intimidated, began by saying, “I guess you want my resignation.” Watson replied, “You can’t be serious. We just spent \$10 million educating you.”

Fortunately, the learning process need not be so expensive. Case studies and post-project reviews like those of Xerox and British Petroleum can be performed with little cost other than managers’ time. Companies can also enlist the help of faculty and students at local colleges or universities; they bring fresh perspectives and view internships and case studies as opportunities to gain experience and increase their own learning. A few companies have established computerized data banks to speed up the learning process. At Paul Revere Life Insurance, management requires all problem-solving teams to complete short registration forms describing their proposed projects if they hope to qualify for the company’s award program. The company then enters the forms into its computer system and can immediately retrieve a listing of other groups of people who have worked or are working on the topic, along with a contact person. Relevant experience is then just a telephone call away.

4. Learning from others

Of course, not all learning comes from reflection and self-analysis. Sometimes the most powerful insights come from looking outside one’s immediate environment to gain a new perspective. Enlightened managers know that even companies in completely different businesses can be fertile sources of ideas and catalysts for creative thinking. At these organizations, enthusiastic borrowing is replacing the “not invented here” syndrome. Milliken calls the process SIS, for “Steal Ideas Shamelessly”; the broader term for it is benchmarking.

Enthusiastic borrowing is replacing the “not invented here” syndrome.

According to one expert, “benchmarking is an ongoing investigation and learning experience that ensures that best industry practices are uncovered, analyzed, adopted, and implemented.”⁷ The greatest benefits come from studying *practices*, the way that work gets done, rather than results, and from involving line managers in the process. Almost anything can be benchmarked. Xerox, the concept’s creator, has applied it to billing, warehousing, and automated manufacturing. Milliken has been even more creative: in an inspired moment, it benchmarked Xerox’s approach to benchmarking.

Unfortunately, there is still considerable confusion about the requirements for successful benchmarking. Benchmarking is not “industrial tourism,” a series of ad hoc visits to companies that have received favorable publicity or won quality awards. Rather, it is a disciplined process that begins with a thorough search to identify best-practice organizations, continues with careful study of one’s own practices and performance, progresses through systematic site visits and interviews, and concludes with an analysis of results, development of recommendations, and implementation. While time-consuming, the process need not be terribly expensive. AT&T’s Benchmarking Group estimates that a moderate-sized project takes four to six months and incurs out-of-pocket costs of \$20,000 (when personnel costs are included, the figure is three to four times higher).

Benchmarking is one way of gaining an outside perspective; another, equally fertile source of ideas is customers. Conversations with customers invariably stimulate learning; they are, after all, experts in what they do. Customers can provide up-to-date product information, competitive comparisons, insights into changing preferences, and immediate feedback about service and patterns of use. And companies need these insights at all levels, from the executive suite to

the shop floor. At Motorola, members of the Operating and Policy Committee, including the CEO, meet personally and on a regular basis with customers. At Worthington Steel, all machine operators make periodic, unescorted trips to customers' factories to discuss their needs.

Customers can provide competitive comparisons and immediate feedback about service. And companies need these insights at all levels, from the executive suite to the shop floor.

Sometimes customers can't articulate their needs or remember even the most recent problems they have had with a product or service. If that's the case, managers must observe them in action. Xerox employs a number of anthropologists at its Palo Alto Research Center to observe users of new document products in their offices. Digital Equipment has developed an interactive process called "contextual inquiry" that is used by software engineers to observe users of new technologies as they go about their work. Milliken has created "first-delivery teams" that accompany the first shipment of all products; team members follow the product through the customer's production process to see how it is used and then develop ideas for further improvement.

Learning organizations cultivate the art of open, attentive listening. Managers must be open to criticism.

Whatever the source of outside ideas, learning will only occur in a receptive environment. Managers can't be defensive and must be open to criticism or bad news. This is a difficult challenge, but it is essential for success. Companies that approach customers assuming that "we must be right, they have to be wrong" or visit other organizations certain that "they can't teach us anything" seldom learn very much. Learning organizations, by contrast, cultivate the art of open, attentive listening.

5. Transferring knowledge

For learning to be more than a local affair, knowledge must spread quickly and efficiently throughout the organization. Ideas carry maximum impact when they are shared broadly rather than held in a few hands. A variety of mechanisms spur this process, including written, oral, and visual reports, site visits and tours, personnel rotation programs, education and training programs, and standardization programs. Each has distinctive strengths and weaknesses.

Reports and tours are by far the most popular mediums. Reports serve many purposes: they summarize findings, provide checklists of dos and don'ts, and describe important processes and events. They cover a multitude of topics, from benchmarking studies to accounting conventions to newly discovered marketing techniques. Today written reports are often supplemented by videotapes, which offer greater immediacy and fidelity.

Tours are an equally popular means of transferring knowledge, especially for large, multidivisional organizations with multiple sites. The most effective tours are tailored to different audiences and needs. To introduce its managers to the distinctive manufacturing practices of New United Motor Manufacturing Inc. (NUMMI), its joint venture with Toyota, General Motors developed a series of specialized tours. Some were geared to upper and middle managers, while others were aimed at lower ranks. Each tour described the policies, practices, and systems that were most relevant to that level of management.

Despite their popularity, reports and tours are relatively cumbersome ways of transferring knowledge. The gritty details that lie behind complex management concepts are difficult to communicate secondhand. Absorbing facts by reading them or seeing them demonstrated is one thing; experiencing them personally is quite another. As a leading cognitive

scientist has observed, “It is very difficult to become knowledgeable in a passive way. Actively experiencing something is considerably more valuable than having it described.”⁸ For this reason, personnel rotation programs are one of the most powerful methods of transferring knowledge.

In many organizations, expertise is held locally: in a particularly skilled computer technician, perhaps, a savvy global brand manager, or a division head with a track record of successful joint ventures. Those in daily contact with these experts benefit enormously from their skills, but their field of influence is relatively narrow. Transferring them to different parts of the organization helps share the wealth. Transfers may be from division to division, department to department, or facility to facility; they may involve senior, middle, or first-level managers. A supervisor experienced in just-in-time production, for example, might move to another factory to apply the methods there, or a successful division manager might transfer to a lagging division to invigorate it with already proven ideas. The CEO of Time Life used the latter approach when he shifted the president of the company’s music division, who had orchestrated several years of rapid growth and high profits through innovative marketing, to the presidency of the book division, where profits were flat because of continued reliance on traditional marketing concepts.

Line to staff transfers are another option. These are most effective when they allow experienced managers to distill what they have learned and diffuse it across the company in the form of new standards, policies, or training programs. Consider how PPG used just such a transfer to advance its human resource practices around the concept of high-commitment work systems. In 1986, PPG constructed a new float-glass plant in Chehalis, Washington; it employed a radically new technology as well as innovations in human resource management that were developed by the plant manager and his staff. All workers were organized into small, self-managing teams with responsibility for work assignments, scheduling, problem solving and improvement, and peer review. After several years running the factory, the plant manager was promoted to director of human resources for the entire glass group. Drawing on his experiences at Chehalis, he developed a training program geared toward first-level supervisors that taught the behaviors needed to manage employees in a participative, self-managing environment.

As the PPG example suggests, education and training programs are powerful tools for transferring knowledge. But for maximum effectiveness, they must be linked explicitly to implementation. All too often, trainers assume that new knowledge will be applied without taking concrete steps to ensure that trainees actually follow through. Seldom do trainers provide opportunities for practice, and few programs consciously promote the application of their teachings after employees have returned to their jobs.

Xerox and GTE are exceptions. As noted earlier, when Xerox introduced problem-solving techniques to its employees in the 1980s, everyone, from the top to the bottom of the organization, was taught in small departmental or divisional groups led by their immediate superior. After an introduction to concepts and techniques, each group applied what they learned to a real-life work problem. In a similar spirit, GTE’s Quality: The Competitive Edge program was offered to teams of business-unit presidents and the managers reporting to them. At the beginning of the 3-day course, each team received a request from a company officer to prepare a complete quality plan for their unit, based on the course concepts, within 60 days. Discussion periods of two to three hours were set aside during the program so that teams could begin working on their plans. After the teams submitted their reports, the company officers studied them, and then the teams implemented them. This GTE program produced dramatic improvements in quality, including a recent semifinalist spot in the Baldrige Awards.

GTE proved knowledge is more likely to be transferred effectively when the right incentives are in place.

The GTE example suggests another important guideline: knowledge is more likely to be transferred effectively when the right incentives are in place. If employees know that their plans will be evaluated and implemented—in other words, that their learning will be applied—progress is far more likely. At most companies, the status quo is well entrenched; only if managers and employees see new ideas as being in their own best interest will they accept them gracefully. AT&T has developed a creative approach that combines strong incentives with information sharing. Called the Chairman's Quality Award (CQA), it is an internal quality competition modeled on the Baldrige prize but with an important twist: awards are given not only for absolute performance (using the same 1,000-point scoring system as Baldrige) but also for improvements in scoring from the previous year. Gold, silver, and bronze Improvement Awards are given to units that have improved their scores 200, 150, and 100 points, respectively. These awards provide the incentive for change. An accompanying Pockets of Excellence program simplifies knowledge transfer. Every year, it identifies every unit within the company that has scored at least 60% of the possible points in each award category and then publicizes the names of these units using written reports and electronic mail.

Measuring Learning

Managers have long known that “if you can't measure it, you can't manage it.” This maxim is as true of learning as it is of any other corporate objective. Traditionally, the solution has been “learning curves” and “manufacturing progress functions.” Both concepts date back to the discovery, during the 1920s and 1930s, that the costs of airframe manufacturing fell predictably with increases in cumulative volume. These increases were viewed as proxies for greater manufacturing knowledge, and most early studies examined their impact on the costs of direct labor. Later studies expanded the focus, looking at total manufacturing costs and the impact of experience in other industries, including shipbuilding, oil refining, and consumer electronics. Typically, learning rates were in the 80% to 85% range (meaning that with a doubling of cumulative production, costs fell to 80% to 85% of their previous level), although there was wide variation.

Firms like the Boston Consulting Group raised these ideas to a higher level in the 1970s. Drawing on the logic of learning curves, they argued that industries as a whole faced “experience curves,” costs and prices that fell by predictable amounts as industries grew and their total production increased. With this observation, consultants suggested, came an iron law of competition. To enjoy the benefits of experience, companies would have to rapidly increase their production ahead of competitors to lower prices and gain market share.

Both learning and experience curves are still widely used, especially in the aerospace, defense, and electronics industries. Boeing, for instance, has established learning curves for every work station in its assembly plant; they assist in monitoring productivity, determining work flows and staffing levels, and setting prices and profit margins on new airplanes. Experience curves are common in semiconductors and consumer electronics, where they are used to forecast industry costs and prices.

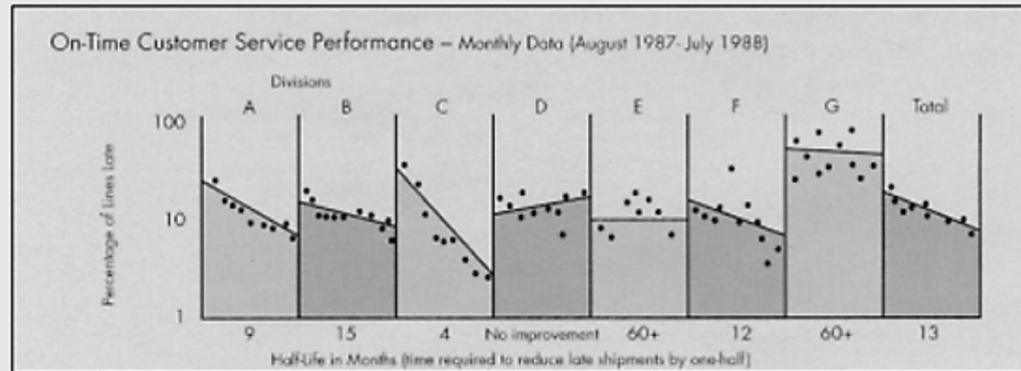
For companies hoping to become learning organizations, however, these measures are incomplete. They focus on only a single measure of output (cost or price) and ignore learning that affects other competitive variables, like quality, delivery, or new product introductions. They suggest only one possible learning driver (total production volumes) and ignore both the possibility of learning in mature industries, where output is flat, and the possibility that learning might be driven by other sources, such as new technology or the challenge posed by competing products. Perhaps most important, they tell us little about the sources of learning or the levers of change.

Another measure has emerged in response to these concerns. Called the “half-life” curve, it was originally developed by Analog Devices, a leading semiconductor manufacturer, as a way of comparing internal improvement rates. A half-life curve measures the time it takes to achieve a 50% improvement in a specified performance measure. When represented graphically, the performance measure (defect rates, on-time delivery, time to market) is plotted on the vertical axis, using a logarithmic scale, and the time scale (days, months, years) is plotted horizontally. Steeper slopes then represent faster learning (see the insert “The Half-Life Curve” for an illustration).

The Half-Life Curve

Analog Devices has used half-life curves to compare the performance of its divisions. Here monthly data on customer service are graphed for seven divisions. Division C is the clear winner: even though it started with a high proportion of late deliveries, its rapid

learning rate led eventually to the best absolute performance. Divisions D, E, and G have been far less successful, with little or no improvement in on-time service over the period.



Source: Ray Stata, "Organizational Learning—The Key to Management Innovation," *Sloan Management Review*, Spring 1989, p. 72.

The Half-Life Curve Analog Devices has used half-life curves to compare the performance of its divisions. Here monthly data on customer service are graphed for seven divisions. Division C is the clear winner: even though it started with a high proportion of late deliveries, its rapid learning rate led eventually to the best absolute performance. Divisions D, E, and G have been far less successful, with little or no improvement in on-time service over the period. Source: Ray Stata, "Organizational Learning—The Key to Management Innovation," *Sloan Management Review*, Spring 1989, p. 72.

The logic is straightforward. Companies, divisions, or departments that take less time to improve must be learning faster than their peers. In the long run, their short learning cycles will translate into superior performance. The 50% target is a measure of convenience; it was derived empirically from studies of successful improvement processes at a wide range of companies. Half-life curves are also flexible. Unlike learning and experience curves, they work on any output measure, and they are not confined to costs or prices. In addition, they are easy to operationalize, they provide a simple measuring stick, and they allow for ready comparison among groups.

Yet even half-life curves have an important weakness: they focus solely on results. Some types of knowledge take years to digest, with few visible changes in performance for long periods. Creating a total quality culture, for instance, or developing new approaches to product development are difficult systemic changes. Because of their long gestation periods, half-life curves or any other measures focused solely on results are unlikely to capture any short-run learning that has occurred. A more comprehensive framework is needed to track progress.

Organizational learning can usually be traced through three overlapping stages. The first step is cognitive. Members of the organization are exposed to new ideas, expand their knowledge, and begin to think differently. The second step is behavioral. Employees begin to internalize new insights and alter their behavior. And the third step is performance improvement, with changes in behavior leading to measurable improvements in results: superior quality, better delivery, increased market share, or other tangible gains. Because cognitive and behavioral changes typically precede improvements in performance, a complete learning audit must include all three.

Surveys, questionnaires, and interviews are useful for this purpose. At the cognitive level, they would focus on attitudes and depth of understanding. Have employees truly understood the meaning of self-direction and teamwork, or are the terms still unclear? At PPG, a team of human resource experts periodically audits every manufacturing plant, including extensive interviews with shop-floor employees, to ensure that the concepts are well understood. Have new approaches to customer service been fully accepted? At its 1989 Worldwide Marketing Managers' Meeting, Ford presented participants with a series of hypothetical situations in which customer complaints were in conflict with short-term dealer or company profit goals and asked how they would respond. Surveys like these are the first step toward identifying changed attitudes and new ways of thinking.

To assess behavioral changes, surveys and questionnaires must be supplemented by direct observation. Here the proof is in the doing, and there is no substitute for seeing employees in action. Domino's Pizza uses "mystery shoppers" to assess managers' commitment to customer service at its individual stores; L.L. Bean places telephone orders with its own operators to assess service levels. Other companies invite outside consultants to visit, attend meetings, observe employees in action, and then report what they have learned. In many ways, this approach mirrors that of examiners for the Baldrige Award, who make several-day site visits to semifinalists to see whether the companies' deeds match the words on their applications.

Finally, a comprehensive learning audit also measures performance. Half-life curves or other performance measures are essential for ensuring that cognitive and behavioral changes have actually produced results. Without them, companies would lack a rationale for investing in learning and the assurance that learning was serving the organization's ends.

First Steps

Learning organizations are not built overnight. Most successful examples are the products of carefully cultivated attitudes, commitments, and management processes that have accrued slowly and steadily over time. Still, some changes can be made immediately. Any company that wishes to become a learning organization can begin by taking a few simple steps.

The first step is to foster an environment that is conducive to learning. There must be time for reflection and analysis, to think about strategic plans, dissect customer needs, assess current work systems, and invent new products. Learning is difficult when employees are harried or rushed; it tends to be driven out by the pressures of the moment. Only if top management explicitly frees up employees' time for the purpose does learning occur with any frequency. That time will be doubly productive if employees possess the skills to use it wisely. Training in brainstorming, problem solving, evaluating experiments, and other core learning skills is therefore essential.

Another powerful lever is to open up boundaries and stimulate the exchange of ideas. Boundaries inhibit the flow of information; they keep individuals and groups isolated and reinforce preconceptions. Opening up boundaries, with conferences, meetings, and project teams, which either cross organizational levels or link the company and its customers and suppliers, ensures a fresh flow of ideas and the chance to consider competing perspectives. General Electric CEO Jack Welch considers this to be such a powerful stimulant of change that he has made "boundarylessness" a cornerstone of the company's strategy for the 1990s.

Once managers have established a more supportive, open environment, they can create learning forums. These are programs or events designed with explicit learning goals in mind, and they can take a variety of forms: strategic reviews, which examine the changing competitive environment and the company's product portfolio, technology, and market positioning; systems audits, which review the health of large, cross-functional processes and delivery systems; internal benchmarking reports, which identify and compare best-in-class activities within the organization; study missions, which are dispatched to leading organizations around the world to better understand their performance and distinctive skills; and jamborees or symposiums, which bring together customers, suppliers, outside experts, or internal groups to share ideas and learn from one another. Each of these activities fosters learning by requiring employees to wrestle with new knowledge and consider its implications. Each can also be tailored to business needs. A consumer goods company, for example, might sponsor a study mission to Europe to learn more about distribution methods within the newly unified Common Market, while a high-technology company might launch a systems audit to review its new product development process.

Together these efforts help to eliminate barriers that impede learning and begin to move learning higher on the organizational agenda. They also suggest a subtle shift in focus, away from continuous improvement and toward a commitment to learning. Coupled with a better understanding of the "three Ms," the meaning, management, and measurement of learning, this shift provides a solid foundation for building learning organizations.

References

1. Peter M. Senge, *The Fifth Discipline* (New York: Doubleday, 1990), p. 1.
2. Ikujiro Nonaka, "The Knowledge-Creating Company," *Harvard Business Review*, November–December 1991, p. 97.
3. Robert Howard, "The CEO as Organizational Architect: An Interview with Xerox's Paul Allaire," *Harvard Business Review*, September–October 1992, p. 106.
4. Modesto A. Maidique and Billie Jo Zirger, "The New Product Learning Cycle," *Research Policy*, Vol. 14, No. 6 (1985), pp. 299, 309.
5. Frank R. Gulliver, "Post-Project Appraisals Pay," *Harvard Business Review*, March–April 1987, p. 128.
6. David Nadler, "Even Failures Can Be Productive," *New York Times*, April 23, 1989, Sec. 3, p. 3.
7. Robert C. Camp, *Benchmarking: The Search for Industry Best Practices that Lead to Superior Performance* (Milwaukee: ASQC Quality Press, 1989), p. 12.
8. Roger Schank, with Peter Childers, *The Creative Attitude* (New York: Macmillan, 1988), p. 9.
9. Ramchandran Jaikumar and Roger Bohn, "The Development of Intelligent Systems for Industrial Use: A Conceptual Framework," *Research on Technological Innovation, Management and Policy*, Vol. 3 (1986), pp. 182–188.

A version of this article appeared in the July–August 1993 issue of *Harvard Business Review*.

David A. Garvin is the C. Roland Christensen Professor at Harvard Business School.

This article is about PERFORMANCE MEASUREMENT

 FOLLOW THIS TOPIC

Related Topics:

CHANGE MANAGEMENT | ORGANIZATIONAL STRUCTURE | MANUFACTURING | KNOWLEDGE MANAGEMENT | STRATEGY EXECUTION

Comments

Leave a Comment

POST

0 COMMENTS

 [JOIN THE CONVERSATION](#)

POSTING GUIDELINES

We hope the conversations that take place on HBR.org will be energetic, constructive, and thought-provoking. To comment, readers must sign in or register. And to ensure the quality of the discussion, our moderating team will review all comments and may edit them for clarity, length, and relevance. Comments that are overly promotional, mean-spirited, or off-topic may be deleted per the moderators' judgment. All postings become the property of Harvard Business Publishing.