2

Organizational Learning Mechanisms

Despite its ready acceptance by researchers, consultants, and managers, the claim that organizations can learn is problematic because it implicitly attributes a human capacity, namely learning, to organizations, which are nonhuman entities. This problem or fallacy of anthropomorphism can be illustrated by comparing Figure 2.1A, which presents a slight adaptation of a well-known model of *how* individuals learn from experience (Kolb, 1984), with Figure 2.1B, which presents a model of organizational learning (Shaw & Perkins, 1992).

The two figures are virtually identical except for "dissemination," which does not fit into the star-shaped configuration, in that it depicts both individual and organizational learning as a sequence of information-processing operations. However, although it is clear how individuals can perform these operations owing to their central nervous system, it is not at all clear how organizations perform them. Furthermore, the awkward position of "dissemination" in Figure 2.1B indicates that nontrivial features of organizational learning may have no analogue in individual learning.

Figure 2.1 Organizational learning

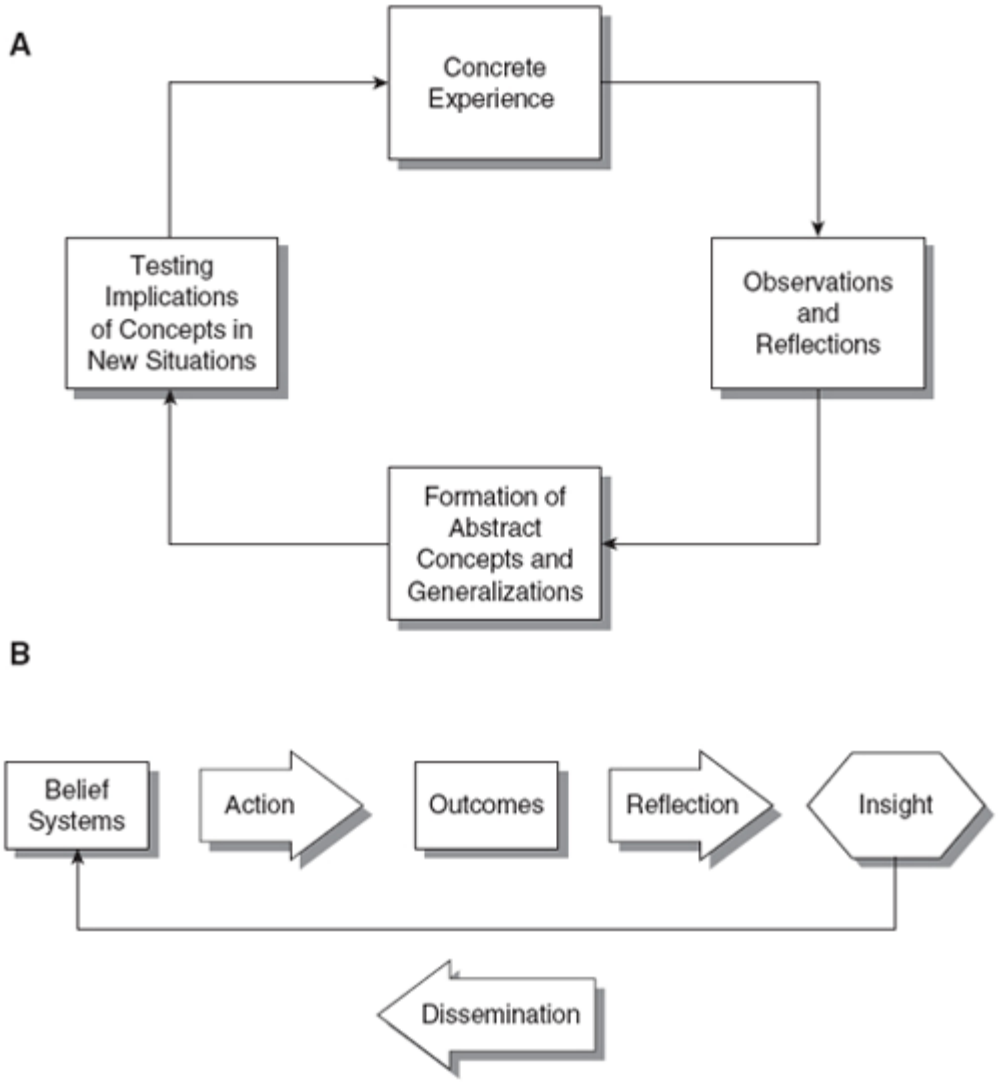

The comparison of Figure 2.1A and Figure 2.1B leads, therefore, to the conclusion that in order to learn, organizations must posses structures analogous to the central nervous system that enables their members (who can learn) to collect, analyze, disseminate, and apply information and knowledge. These structures are organizational learning mechanisms (OLMs), like after-action and postproject reviews, communities of practice, and benchmarking teams. In our approach, the study of organizational learning begins with the identification of organizational learning mechanisms and the exploration of their natures.

The concept of OLMs is useful for several reasons. First, it highlights the similarities and differences between individual and organizational learning. The two are similar inasmuch as they both involve the processing of information and knowledge (see Chapter 7 on dissemination). They are dissimilar inasmuch as both the nature of the systems with which information is processed and the nature of these processes in each one of them are different. In the individual case, information is processed by the central nervous system, which produces individual level output such as changes in mental models and individual behavior. In the organizational case, information is processed by groups of people interacting within organizational learning mechanisms, which produce organizational level output in the form of

changes in shared mental models, formal procedures, and informal norms. As a corollary, while individual learning is a mental or cognitive process, organizational learning is primarily a social process.

The second contribution of OLMs is that they help to demystify organizational learning by providing a nonmetaphorical way of showing how organizations learn. Through the operations of OLMs, information and knowledge relevant to the organization's survival and prosperity are collected, analyzed, disseminated, and stored for future use. OLMs also provide a means of demonstrating how individual level learning becomes organizational level learning. Through the interactions in OLMs, individual knowledge and learning are transformed into changes in organizational routines, standard operating procedures, shared beliefs, and informal norms.

Organizational learning mechanisms have been discussed by other researchers without necessarily using this terminology (Baird et al., 1997; Carroll, 1995; Cheney, 1998; DiBella et al., 1996; Dodgson, 1993; Gulliver, 1987; Shani & Docherty, 2003; Shani & Mitki, 2000). The purpose of this chapter is to demonstrate the conceptual and practical usefulness of the concept of OLMs by presenting a basic typology that organizes every OLM we have encountered so far into one of five basic categories (Figure 2.1, Row 4). In doing so, we will illustrate typical OLMs and examples of best practice within each category. We hope to demonstrate the conceptual power of OLMs by their ability to integrate a large and fairly disparate literature within a single conceptual framework.

We also hope to illustrate the practical usefulness of the OLM concept in two ways. First, analyzing organizational learning in a particular organization can begin by identifying and characterizing the OLMs that it operates. Second, the first step in initiating organizational learning in a system is to design OLMs that are suitable for its particular circumstances. The basic typology and detailed specific examples presented in this chapter should be helpful in both respects. It may also be useful to refer back to Figure 2.1 at the beginning of each subsection to keep track of how the discussion proceeds as we move from one type of OLM to another.

❖ TYPES OF ORGANIZATIONAL LEARNING MECHANISMS

Given the wide variety of OLMs, it is useful to have a classification scheme that organizes them and highlights their essential characteristics. We propose a two-dimensional scheme for OLMs derived from two basic questions regarding the nature of learning: Who carries out the learning? When/Where does learning take place relative to the task itself? The first question regards the “agents” of learning; that is, who collects information about a particular task or problem, analyzes this information, draws conclusions, and disseminates the products of learning? Agency entails the extent that learning is carried out by the same people who perform the task. “Internal” agents of learning are those people who engage in both task performance and learning about task performance. “External” agents are people whose role is to engage and promote learning about tasks performed by others.

The second question regards the temporal and spatial relations between task performance and the act of learning. This question concerns itself with the extent to which learning about the task (collecting, analyzing, and disseminating information) takes place at the same time and place as the work itself. “Online” learning takes place in very close proximity to task performance itself; “off-line” learning takes place at a distinctly different time, and often a different place, than the work itself. These two characteristics yield four basic types of organizational learning mechanisms: online/internal agency, online/external agency, off-line/internal agency, off-line/external agency. This simple two-way classification is exhaustive:

any organizational learning mechanism—and hence any type of systematic organizational learning activity—can be classified as one of these types. In the remainder of the chapter, we will discuss specific OLMs in each category based on their prominence in the literature:

- After-action reviews, postproject reviews, and communities of practice (off-line/internal OLMs)
- Online experimentation and online debriefing (online/internal OLMs)
- Postproject assessment units and scenario planning units (off-line/external OLMs)
- Coaching networks and peer assists (online/external OLMs)

The four basic OLM types formed by the two dimensions are primarily concerned with learning (knowledge creation) and to a lesser degree with knowledge dissemination. A fifth type of external organizational learning mechanism is specifically dedicated to knowledge dissemination. These mechanisms are discussed in Chapter 7 and included in Figure 2.1 only for the sake of closure.

❖ OFF-LINE/INTERNAL OLMs

After-Action Reviews (AAR) and Postproject Reviews

These OLMs, which are arguably the most thoroughly discussed in the literature, are predicated on two principles. The first principle is that organizations accumulate valuable knowledge that is directly relevant to their success in the course of their operation, so they benefit from critical reflection on past experience by their individual members, units, or the entire organization. The second principle is that reflection is best carried out by the same individuals who participated in the action or project because they are the ones responsible for its outcomes and for the implementation of lessons learned. There are potential dangers to this approach because individuals may be subject to pressures to cover up errors and shirk responsibility for failure. However, choosing this OLM assumes that the advantages outweigh the disadvantages.

After-action reviews are frequently practiced in the military following training exercises and combat operations. One best practice example of this OLM—postflight reviews in a combat squadron of the Israel Defense Force Air Force—is analyzed in detail in Chapter 8. Chapter 10 presents a general model of high-quality organizational learning that we developed by comparing successful and unsuccessful postaccident reviews in an elite combat unit of the Israel Defense Force. The U.S. Army provides another best practice example of systematic after-action reviews. Margaret Wheatley observed the U.S. Army's AARs, and her vivid description captures the difficult requirements for openness that require successful AARs to be embedded in a culture that promotes transparency, inquiry integrity, and issue orientation (see Chapters 3, 8, and 10 for further discussion of this point):

> During an AAR soldiers and their commanders gather to probe in detail: “What went on…” “What actions occurred…” “What actions should have occurred…” Who did what…” Every participant in these AARs is under a microscope, being called to accountability by their colleagues for their decisions and actions during the time under review. Behaviors of the leader, individuals, and the group are all analyzed. They are equal, engaged in discovery learning. Soldiers describe these sessions as brutal, sweat inducing—and absolutely necessary for their learning. The process has become so ingrained that during the Gulf War AARs would be held spontaneously in the rear of a

> truck—called by anyone independent of rank. These AARs have fostered an expectation that behaviors and decisions will be reviewed, and that everyone will benefit from the learning, no matter how difficult it is to hear at the time. (Wheatley, 1994, p. 52)

Boston University's researcher Lloyd Baird and his associates identified the basic characteristics of AARs as *focused* on few critical issues; done *immediately after* the action; *inclusive* of all those who took part in the action; following a *structured* process; and *leading back to action* as soon as possible (Baird et al., 1997.) The purpose of focus is to enable participants to identify a few key factors that are important to improve performance in less than an hour and get back to action. Proximity to the event helps participants recall what happened most accurately. Finally, inclusiveness helps to gain all relevant perspectives on what happened (and, in our opinion, to maximize the understanding and commitment of all relevant parties to the implementation of their lessons learned).

Baird and his associates also proposed a simple five-step procedure for effective AARs, which we have slightly revised as follows:

1. Establish what was the intent or purpose of action.
2. Establish what happened exactly—why, how, and what were the results?—by asking participants to reconstruct what happened chronologically.
3. Determine lessons learned: What new facts emerged in the previous two steps, and what can be done on this basis to improve action next time around?
4. Identify which other parties should know the lessons learned and how they are going to be told in order to ensure their implementation and to improve the performance of other units in the organization.
5. Take action (lessons learned that do not lead to action do not matter).

Finally, Baird and his associates (1997) recommend that AAR facilitators should (a) "follow the rule of objectivity" by asking participants to report the facts of the events and to separate these facts from their interpretations; (b) "balance inquiry and advocacy" by asking participants to focus on other participants' factual reports and explanations as much as on their own explanations; (c) "climb the ladder of inference" by encouraging participants to proceed systematically from direct observation through their interpretation on the basis of past experience to final conclusions; and (d) "ready, fire, aim, aim"; that is, encourage participants to take action and complete their learning as they go, in the manner of the smart bomb that corrects its course by tracking its target as it goes.

Despite their ubiquity in the military, AARs and postproject reviews are not confined by any means to it. Microsoft, where post-project reviews are called "postmortems," provides another best practice example of how this OLM is carried out. In a high-tech environment, people's familiarity and comfort with electronic media allows the OLM to engage several layers of the organization in the process, conduct the reviews in a distributed rather than face-to-face fashion, and retain their effectiveness by holding on to norms of transparency, inquiry, integrity, and issue orientation (see Chapter 3 for a discussion of the cultural norms that generate productive learning):

> Since the late 1980s, between half and two thirds of all Microsoft projects have written postmortem reports and most other projects have held postmortem discussion sessions. The postmortem documents are surprisingly candid in their self-criticism, especially because they are circulated to the highest levels of the company…. Groups generally take three to six months to put a postmortem document together. The documents have ranged from under 10 pages to more than 100 pages, and have tended to grow in length…. The

> most common format is to discuss what worked well in the last project, what did not work well, and what the group should do to improve in the next project…. The functional managers usually prepare an initial draft and then circulate this via e-mail to the team members, who send in their comments. The authors collate these and create the final draft, which then goes out to team members as well as senior executives and directors of product development, and testing. The functional groups, and sometimes an entire project, will then meet to discuss the postmortem findings. Some groups… have also gotten into the habit of holding postmortem meetings at every milestone to make mid-course corrections, review feature lists, and rebalance schedule…. The Excel team took the lead in finding solutions for relatively large projects. (Cusumano & Selby, 1995, pp. 331–332)

Communities of Practice

Communities of practice are groups whose members meet regularly to share knowledge and learn together in areas of joint concern (Lesser & Storck, 2001). Thus, communities of practice are similar to AARs and postproject reviews in that both are predicated on the assumption that organizations can benefit by sharing knowledge that accumulates "inside the heads" of their own members. Different from the latter, they are not tied to or focused on specific shared experiences but rather enable members with different levels of expertise to share information and knowledge and reflect on problems of common interest.

It is possible to get a sense of how communities of practice operate from the testimony of a member of Microsoft Excel's testing team:

> On the testing level, I speak quite frequently with the Word testing managers…. We have testers hooked up who are working on similar features across their groups, so they share ideas and information. We meet once a month right now for lunch for two hours—the test managers from Word, Excel and Project. We talk about "what are we facing?" "How did you solve this problem?" "I am thinking about this issue, what did you guys do?"…We meet monthly with all the test managers within the company, within the Worldwide Product Group…. We do a presentation, and we all share what our groups are doing. (Cusumano & Selby, 1995, p. 342)

Participation in a community of practice differs from both working within a team and from participation in training. Whereas communities of practice are voluntary and focus on improving professional skills, teams focus on performing job-related tasks and are assigned by the organization. And whereas training activities (workshops, conferences, and courses) typically last several days or weeks, communities of practice may last for several years. Most important, communities of practice and training differ in terms of the nature of the knowledge that their participants acquire. Courses and workshops provide explicit textbook-type knowledge that can be expressed verbally or numerically and that usually seek to apply across contexts and situations. The type of knowledge delivered by communities of practice is technically known as "tacit" and "situated."

Tacit knowledge refers to "the ability to do things without being able to explain them completely and also the inability to learn to do them from a theoretical understanding of the task" (Orr, 1990, p. 170.) For example, knowing how to talk, walk, or ride a bicycle is tacit knowledge. It cannot be learned by reading an instruction manual but can be acquired by watching proficient performers and by receiving corrective feedback from them. *Situated knowledge* (Lave, 1991) is embedded in the particular context of practice in which it arises or applies. For example, it means not just knowing how to drive, but how to drive this or that

vehicle, in this or that terrain, or for this or that purpose. Another example of situated knowledge is knowing the appropriate actions in specific situations that arise as part of fulfilling a particular job in a particular organization (see Chapter 7 for additional discussion of the differences between tacit and explicit knowledge and its importance for organizational learning).

Communities of practice are self-organized, which means that they cannot be mandated by the parent organization. What organizations can do, however, is encourage their formation and support their functioning. Specifically, organizations can provide communities with infrastructure such as official sponsors and support teams, which help them operate. Management can support them by assessing their value to the organization through nontraditional methods such as members' stories about how communities of practice contributed to learning and improvements (Orr, 1990). In addition, the parent organization can identify potential communities and help them to organize by assisting people who wish to start a community to reach prospective members.

British Petroleum, for example, uses a voluntary corporate Yellow Pages system, dubbed "Connect," as the platform for making networks visible. Connect serves as a directory to BP's knowledge workers and associated networks. Originally conceived as a way for technical staff to articulate their capabilities, BP's Connect system has grown to include the Web-based personal profiles of more than 18,000 knowledge workers and more than 250 networks. Knowledge workers use Connect to select networks they wish to join, to locate individuals with common skills and interests, and as an aid in matching people with needed skills and experience to particular project needs. Connect also provides a way for all of the networks in the corporation to be recognized and to operate transparently—in such a way that everybody can see what everybody is doing (Barrow, 2001; Prokesch, 1997).

❖ ONLINE/INTERNAL OLMs

Online/internal learning means that working and learning are fused together: Task performance becomes an organizational learning mechanism. This form of learning materializes when work is accompanied by certain practices that produce changes in organizational routines, standard operating procedures, or norms of behavior. Basically, it is the organizational level analog of "reflection-in-action," Donald Schön's term for the skill of proficient practitioners who combine action with critical reflection on that action (Schön, 1983). Reflection-in-action accounts for the artistry of professional performance in the short term and high levels of performance in the long term. Two practices of internal/online organizational learning are reported in the literature: online experimentation and online debriefing.

Online Experimentation

This form of internal/online organizational learning corresponds to "on-the-spot" experiments, Schön's term for testing hidden assumptions unearthed by reflection-in-action or for testing a new course of action generated this way. These experiments are usually carried out by acting on the assumptions or implementing the new courses of action and observing the compatibility of their outcomes with the practitioner's expectations. The organizational level analogue of on-the-spot experiments was developed in Chaparral Steel (see Chapter 11 for an in-depth analysis of online/internal learning) where responsibility for R & D and work improvement is assigned to line workers so that "many creative simulations are conducted

right on the production line" by the production workers themselves (Leonard-Barton, 1992, p. 31).

Online Debriefing

Online debriefing is a form of deliberate practice that denotes the mindful way in which experts practice and constantly hone their skills. It was coined by Anders Ericsson to describe the key he and his colleagues discovered for achieving high-level performance (Ericsson, Krampe, & Tesch-Romer, 1993). Online debriefing is similar to on-the-spot experimentation in the requirement for mindfulness but differs in that it does not involve testing in any formal or informal sense. Rayner provided a vivid description of the online debriefing that emerged at Globe Metallurgical, Inc., a steel mill, when its unionized workers went on strike:

> As the union workers left the plant, about 35 salaried workers and 10 company managers stepped in to take over operation of two of the five furnaces…. I [the General Manager, Sims] was assigned to work on the maintenance crew, the dirtiest job in the whole plant. I still don't know who made the assignments…. The strike was a time of great stress but also a time of great progress. We experimented with everything…. A few weeks after management took over operating the plant, output actually improved by 20%…. We were operating in a very fast, continuous improvement mode. Every day people would suggest ways to improve the operation of the furnaces or the additive process or the way we transported material around the plant. I kept a pocket notebook, and if I saw something I'd note it down and discuss it with the team over coffee or during meals. I filled a notebook every day…. As we made more changes and as we settled into the routine of running the plant, we didn't need first-line supervisors. We could produce the product more effectively if everyone just worked together cooperatively—welders, crane operators, furnace operators, forklift drivers, stokers, furnace tapers, and taper assistants. (Rayner, 1993, pp. 287–289)

The example of the emergence of online/internal learning at Globe Metallurgical highlights a conceptual problem that arises when the distinction between working and learning is blurred. It is fair to assume that people usually learn something when they work, most certainly if they engage in reflection-in-action. Does that mean that at all learning on the job is a form of organizational learning and that all work is essentially an online/internal agent OLM? We do not think so. Rather, we suggest that only reflection-in-action (e.g., online experimentation or deliberate practice) that leads to organizational or unit level changes constitute online/internal OLMs.

❖ OFF-LINE/EXTERNAL OLMs

The basic assumption that differentiates external from internal, off-line OLMs is that learning is best performed by experts. Usually, these experts are assigned to this task on a full-time basis, possess specialized analytic skills, and work in centralized units that serve as organizationwide repositories of knowledge and as knowledge dissemination centers. Three best practice OLMs that fall under this heading are the U.S. Army Center for Army Lessons Learned (CALL) (Baird et al., 1997), British Petroleum's Post-Project Assessment Unit (Gulliver, 1987), and Shell's Strategic Scenario Planning Unit (De Geus, 1988; van der Heijden, 1996).

Postproject Assessment Units

The Center for Army Lessons Learned, or CALL, was established in 1973 with a mandate to observe Army training activities and identify lessons learned for improving future training activities. In 1992–1993, this mandate was expanded to include lessons learned for future operational activities. CALL teams are assigned to observe operational activities such as the U.S. Army's peace missions in Ethiopia, Somalia, and Rwanda. Boston University researchers Baird and colleagues' (1997) description of CALL reveals that the center operates in three capacities: developing the Army's simple format of after-action reviews and training units in its implementation; observing training activities, collecting lessons learned, and transforming then into training materials such as combat scenarios; and collecting lessons learned from observation of actual operations and transforming these into training materials tailor-made for troops going into action. In short, CALL is both a collector and repository of lessons learned for the entire Army and a developer of forcewide capacity for the collection of lessons learned independently of its own activities:

> Observers from CALL were assigned to troops going into Haiti [in order to capture] lessons-learned as the troops deployed, quickly identifying critical knowledge and skills needed, and immediately imbedding them in the training program of troops to follow. An example of CALL's success is the transition of troops from the Schoenfeld Barracks in Hawaii to replace the 10th Mountain Division. Immediately upon receiving notification they would be the replacements, troops at Schoenfeld Barracks began training using 26 scenarios developed by CALL from observing the 10th Mountain Division. The 26 scenarios included situations faced by the first troops in and suggested best solutions, complete with video footage of the actual events, virtual simulations, and scripted responses. (Baird et al., 1997, p. 387)

Whereas CALL is designed to serve the entire U.S. Army with a strong focus on helping its operational units, British Petroleum's Post-Project Assessment (PPA) Unit is designed is to help the energy giant's corporate management (Gulliver, 1987.) The unit is located in the corporate headquarters of this conglomerate, which consists of 11 independent business units, each with its own board of directors and chief executive officer. The mission of the unit is to study selected projects, identify generally applicable mistakes and best practices, disseminate this knowledge, and help implement it throughout the corporation. In 1987, the unit was staffed by a manager and four assistants, who reported directly to British Petroleum's board of directors. British Petroleum's board of directors approves the projects selected by PPA for investigation on the basis of potential interest for the entire organization. Each project is studied by a team of two or three PPA staff members from the time it was conceived through its various stages—proposal, construction, up to the first 2 years or so of operation. Six projects are appraised each year.

The duration of the appraisal of a large project is about 6 months. The team begins by spending 2 months or so studying project and corporate files (e.g., accounting, legal, or planning) to familiarize itself with the background information such as the economic climate at the time of the project's inception, its objectives and planned timetable, methods of operation, and identity of the contractors. Next, pairs of team members interview on average some 40 people involved in the project. In addition to factual material, the interviewers collect impressions regarding the psychology of the project members and managers. By piecing together the information collected from various sources, the team creates a broad integrative report regarding the important factors that contributed to the project's problems or its success.

Before submitting reports to the board of directors, PPA allows the managers of the appraised project to inspect it and dispute its fairness and validity.

The fact that the PPA unit enjoys the cooperation of the appraised units and does not encounter their resistance is amazing, given that project managers have been reprimanded for problems unearthed by appraisals of their projects. This can be attributed to three factors. First, British Petroleum has gained the commitment of its managers and workers, who are genuinely interested in helping the organization correct mistakes and benefit from successes. Second, the postproject appraisals are perceived as a useful source of knowledge to which managers themselves can turn when seeking help in solving their own problems. Finally, the corporation has persuaded its members that it values learning by instituting a variety of OLMs in addition to the PPA and the communities of practice, which we reviewed earlier in this chapter.

Scenario-Planning Units

There was a time when planning, which is future oriented, was associated with forecasting rather than learning, and particularly learning from past experience. This perception has changed, thanks to the work of a group of scenario planners in another energy giant, Royal Dutch Shell. Both British Petroleum and Royal Dutch Shell attribute their long-term survival at the top of the world's oil industry to their ability to learn and adapt in a market that experiences unanticipated shortages followed by periods of glut. Within this common denominator, the two arch rivals adopted very different approaches to organizational learning. British Petroleum opted for learning from past experience, to which end it developed a wide array of OLMs such as the Post-Project Appraisal Unit, as well as a network of communities of practice supported by state-of-the-art information technologies.

In contrast, Royal Dutch Shell opted for learning by planning, a notion advanced by Arie de Geus, former head of planning for Royal Dutch Shell's group of companies (De Geus, 1988). De Geus defined institutional (organizational) learning as follows:

> The ability of a company's senior managers to absorb what is going on in the business environment and to act on that information with appropriate business moves. [This ability is achieved through a] process whereby management teams change their shared mental models of their company, their markets, and their competitors. For this reason we [Shell's department of corporate planning] think of planning as learning and of corporate planning as institutional learning. (De Geus, 1988, p. 70)

The notion of planning as learning is very different from the ordinary notion of planning, which is deciding on future courses of action. As we have already noted, planning is traditionally associated with forecasting—a plan's effectiveness is clearly contingent on accurate prediction of the future for which the plan is intended. During the 1950s and 1960s, corporations began to rely on sophisticated quantitative forecasting methodologies, which performed well in the relatively stable business environments of that period.

These methods lost their luster as the increased pace, competitiveness, and unexpected twists and turns of the following decades made prediction of the future by extrapolating from the past less and less viable. In response to this adversity, corporate planners in Shell and elsewhere replaced forecasting by *scenario building*, the principal tool for planning as learning. The purpose of scenario building is not to predict the future but to change manager's shared mental models (which De Geus calls “world views”), thereby sensitizing them to a variety of potential eventualities. This requires thorough information search and analysis by a group of specialists:

> In contrast to strategic plans which historically have been built on line and range forecasts associated with probabilities, scenarios present several starkly contrasting futures, none of which is "right" or more likely than the next, and each of which is plausible. The theory… [of] scenario building is that, if executives are aware of and at least modestly prepared for several possible outcomes, they will be better prepared to adjust if the world takes an unusual turn.
>
> [We at Shell's planning department see] our task as producing a documented view of the future business environment five or ten years ahead. Our real target [is]…to design scenarios so that managers would question their own model of reality and change it when necessary; so as to come up with strategic insights beyond their minds' previous reach.
>
> But exposing and invalidating an obsolete worldview is not where scenario analysis stops. Constructing a new model is the most important job and is the responsibility of the managers themselves. The planners' job is to engage the decision makers' interest and participation in this reconstruction. We listen carefully to their needs and give them the highest quality materials in making decisions. The planners will succeed, however, only if they can securely link the new realities of the outside world—the unfolding business environment—to the managers' microcosm. Good scenarios supply this vital bridge; they must encompass both mangers' concerns and external reality. Otherwise no one will bother to cross the bridge. (Wack, 1985, pp. 80–87)

Numerous success stories (see Fahey & Randall, 1998; Schoemaker & Schuurmans, 2003) show that Shell's early success with scenario planning can be duplicated by other organizations. This is not to say that scenario planning is fail-safe, shielding users from nasty surprises. Shell itself learned this lesson the hard way twice in 1995. First, it insisted on dumping a redundant oil rig, Brent Spar, into the North Sea, only to retract 3 years later after a costly consumer boycott and occupation by a Green Peace team. Then later that year, the Nigerian poet and environmental activist Ken Saro-Wiwa and nine codefendants were hanged by the Nigerian government for protesting the destruction of native land by Shell and other oil companies, generating another international storm of protest and cries to boycott the company (Elkington & Trisoglio, 1996).

The common denominator of learning from after-action or post-project reviews and learning from projections of potential futures is twofold. First, in both cases, learners answer the same questions: How do I understand these events and what implications do they have for my own, my unit's, or my organization's future actions? Secondly, successful learning involves challenging current assumptions and beliefs about the world that drive decision making and action. This challenges results in the development of new more appropriate ways of thinking, which have been referred to as "theories of action" (Argyris & Schön, 1974) or "mental models" (Senge, 1990). The learning processes in post-project reviews and scenario planning are very similar, although one deals with reconstructed past events and the other with plausibly constructed future events.

❖ ONLINE/EXTERNAL OLMs

Online/external OLMs are designed to link organization members with persons with relevant knowledge or expertise who can help them learn or apply this knowledge online. They are exemplified by the extensive analytical staff employed by American football teams who, although not participating in the game itself (carrying out the central task), work with those who do (the players). These analysts actively observe the game and pass information and

strategy to the head coach, who combines it with his own online analyses and then calls in plays to the team itself. Indeed, coaching is the quintessential online/external OLM and is often referred to as "coaching networks" or "peer assists."

Coaching Networks

"Coaching is working with others, in one-on-one relationships to help them achieve breakthroughs in knowledge, work, or thinking" (Bowerman & Collins, 1999, p. 203). These authors report that during the 1990s several Canadian organizations established coaching networks in which "individuals work with each other as coaches and performers to resolve problems that they define for themselves, and in which no immediate solution is readily visible or available" (p. 203).

Peer Assists

Talisayon (2001) described peer assists in the following way:

> Horizontal face-to-face transfer of tacit knowledge across equals, … combining knowledge about what works or what works well, with the knowledge of the assistee about specific local conditions and needs…. The process works best as a horizontal (no pulling of rank) collaborative process that is based on mutual respect and associated "people skills."… The aim of peer assist is for the assistee to be enabled to perform an action better or to achieve a desired result. It is to bring collective knowledge to bear efficiently and effectively toward this end…. The peer assist process takes place during performance of an action, and learning takes place while assister and assistee perform the details of the task together. (pp. 1–2)

Coaching networks and peer assists are very similar in their method of operation and the assumption that organizations have within them useful resources of knowledge that can be called on to help members learn and tackle difficult problems. The difference between the two is that coaching networks assume that knowledge is found only in a select group of experts (who are called on to serve as coaches), whereas peer assists further assume that useful knowledge is distributed throughout the whole of the organization.

Dixon (1999) provided a detailed description of how peer assists work at British Petroleum:

> The "assisters" are not corporate staff, nor are they in any hierarchical or reporting relationship to those who are asking for assistance. Rather, they are peers who, in the coming months, are themselves likely to be asking others to give them the same kind of assistance. These assisters travel to the site of the team that is requesting the assistance and work with that team on a specific issue the team is facing. They do not arrive with a "dog and pony" show; rather the focus of the day-long meeting is on the specific objectives that the team making the request has laid out. (p. 214)

Several aspects of coaching networks as exercised in British Petroleum are worth noting. First, the relationship between the person or group who does the coaching and the person or group who is being coached is determined ad hoc; that is, peer assists are predicated on the assumption that every organization member is potentially a source of valuable knowledge (a practice that is also consistent with the value of issue orientation). Secondly, peer assists will not work unless organization members will perceive coaching as equally important to doing

their own work and without management's support, namely, performance evaluation and remuneration policies that are consistent with this perception.

❖ MECHANISMS, TECHNOLOGIES, AND CHOOSING THE "RIGHT" OLM

Having surveyed a representative sample of the variety of mechanisms that are discussed in the literature, it is important to draw a distinction between *organizational learning mechanisms* and organizational learning (and knowledge management) methodologies and technologies. *Organizational learning methodologies* include learning by joint ventures or the hiring of key personnel, training, reverse engineering, imitating inward investors, best practice transfer, and learning from customers and users. *Organizational learning technologies* include the intranet, internal yellow pages, and Lotus Notes groupware. Organizational learning methodologies and technologies are not themselves OLMs, but they are important for their smooth operation. A particular methodology or information technology (IT) qualifies as an *element* of an OLM once it is systematically employed by a team (or teams) of organizational members for the purpose of learning. Similar to organizational learning in general, OLMs are social entities, as distinct from pieces of hardware or software.

The distinction between OLMs and organizational learning methodologies and technologies is particularly important for selecting a strategy for introducing organizational learning into an organization. Here we just pointed to a relevant distinction between two strategies of knowledge strategies. The first is a *codification* strategy, where knowledge is carefully codified and stored in databases, and the second is a *personalization* strategy, in which knowledge is closely tied to the person who developed it and is shared mainly through person-to-person contacts. The same distinction has been drawn between an *IT* knowledge management track or focus, which "views knowledge as 'objects' to be documented, classified, stored, retrieved, analyzed and otherwise manipulated for useful applications," and a *people* knowledge management track or focus, which "views knowledge as primarily tacit, largely embodied in the skill of experts, embedded in processes intimately linked with people, and often difficult to codify" (Hansen, Nohria, & Tierny, 1999, p. 107).

Each of these strategies or approaches has its advantages and disadvantages. Some of these differences can be traced to the different capabilities that people and IT systems bring to organizational learning and knowledge management.

When we seek to understand knowledge, to interpret it within a broader context, to combine it with other types of information, or to synthesize various unstructured forms of knowledge, humans are the recommended tool. These are the types of knowledge tasks at which we excel, and we should be employed for these purposes. Computers and communications systems, on the other hand, are good at different types of things. For the capture, transformation, and distribution of highly structured knowledge that changes rapidly, computers are more capable than people. They are increasingly useful—though still a bit awkward—for performing these same tasks on less structured textual and visual knowledge. But it is still the case that most people don't turn to computers when they want a rich picture of what is going on in a particular knowledge domain. Given this mixture of skills, we need to construct hybrid knowledge management environments in which we use both humans and people in complementary ways (Talisayon, 2001, p. 1).

One of the disadvantages of the IT/codification approach is that it leads people to focus on organization learning technologies or practices instead of on OLMs and the other facets of organizational learning. Instituting organizational learning in a particular situation requires not merely the importation of the methodologies and technologies but setting up the mechanisms

that enable members to use them as aids for reflection on behalf of the organization. Simply installing an IT system, or even starting an organizational learning or knowledge management initiative by installing one, is a misguided approach to organizational learning. Rather, organizational learning begins with setting up OLMs as well as instituting polices that foster a supportive psychological climate and cultural norms.

What guidelines are there for "choosing the right OLM for my organization?" Our own fundamental position on this subject can be summarized in terms of four principles:

1. OLMs can be both formal and informal organizational entities. Communities of practice, for example, are often informal, though managers who are cognizant of their usefulness can encourage and support their formation, as we have noted in the discussion of this OLM. The common denominator of both formal and informal OLMs is that they are enduring and, in that sense, institutionalized features of the organization.

2. There are no hard and fast rules that match OLM of type X with organizational contingency Y. Two ways to begin the design and systematic implementation of OLMs in a particular organization are the design approach and the replication approach. The design approach consists of considering the four generic types and specific best practice exemplars presented in this chapter and selecting the one—or some variant of one—that best suits the needs and circumstances of the particular organization. The replication approach consists of surveying the OLMs that are already in place in the organization and selecting those, or variants of those, that can be replicated in other units of the organization. For example, Collison and Parcell (1998) describe how