PART I

# THE THEORETICAL LANDSCAPE

# 1

# Toward a New Framework for Understanding Expertise

Studies of expertise commonly start out with extraordinary performances, such as a championship-level chess game. These studies typically conclude that the most important factor behind extraordinary skill is huge amounts of rather repetitive practice. In other words, standard studies of expertise seek mundane processes and explanations behind extraordinary performances.

In this book I move in the opposite direction. I start out with mundane performances in health care workplaces. Behind these mundane work activities I uncover extraordinary potentials and processes that make the continued flow and ongoing transformation of expert work possible.

In other words, this book is not about universal cognitive mechanisms supposedly found in the minds of such prototypical lonesome experts as master chess players or physicists solving well-constrained mathematical problems. This book looks at expertise as everyday work. Such work is carried out by mixed groups and communities of people in conditions where disruptions and unexpected events are the rule rather than the exception. This means, among other things, that expertise is not limited to professionals who have received extensive formal training in their respective fields.

Two classes of mundane events are becoming increasingly pervasive and "normal," yet also increasingly difficult to deal with for traditional studies of expertise. These events are *disturbances* or *breakdowns* on the one hand and *rapid transformations* in the contents of work, technologies, and organizational patterns on the other hand. The two are interconnected. The introduction of novel tasks, technologies, and organizational patterns often increases the likelihood of disturbances and

breakdowns – and recurring disturbances often force practitioners and their management to seek new socio-technical solutions and ways to organize work (Hirschhorn, 1984; Perrow, 1984; Zuboff, 1988). These events make it difficult to build expertise on huge amounts of repetitive practice in relatively stable conditions. The conditions do not remain stable. Experts must face, diagnose, and resolve novel situations for which they often have little or no directly applicable prior practice.

These factors create situations in which employees at all levels of the hierarchy, and increasingly also their clients, face tasks that they find impossible to solve. There is something curious about this impossibility. Each individual, including highly educated professionals and managers, may testify that the situation was clearly beyond his or her control. Yet, most of those situations are somehow resolved and the work goes on. Moreover, often none of the persons involved can quite reconstruct or fully understand what actually happened and how the solution was found. In other words, people at work somehow go beyond their own limitations all the time. What makes this possible is a question I try to answer in this book.

## Traditional Approaches to Expertise

During the past few decades, the cognitive foundations of expertise have been established as a central research theme for cognitive science and artificial intelligence. Despite – and partly because of – important achievements in these fields, our understanding of expert thinking and its formation at work is ready for a major transformation.

There is a pervasive dualism in Western conceptions of human cognition. The dualism is expressed in a number of related versions: analytical vs. intuitive; explicit vs. tacit; scientific vs. experiential; paradigmatic vs. narrative, and so on. Collins (1990, p. 4) characterized the two poles as "algorithmic" and "enculturational" and observed: "We can contrast two models of learning: an 'algorithmic model,' in which knowledge is clearly statable and transferable in something like the form of a recipe, and an 'enculturational model,' where the process has more to do with unconscious social contagion."

In studies of expertise, the algorithmic or human information-processing approach was launched by Herbert Simon and his colleagues in studies of playing chess and solving physics problems (Newell & Simon, 1972; Chase & Simon, 1973; Simon & Simon, 1978). Representative collections of research continuing and expanding on this tradition include

*The Nature of Expertise*, edited by Chi, Glaser, and Farr (1988); *Toward a General Theory of Expertise*, edited by Ericsson and Smith (1991); *The Psychology of Expertise*, edited by Hoffman (1992); *Expertise and Technology*, edited by Hoc, Cacciabue, and Hollnagel (1995); *The Road to Excellence*, edited by Ericsson (1996); *Expertise in Context*, edited by Feltovich, Ford, and Hoffman (1997); *The Cambridge Handbook of Expertise and Expert Performance*, edited by Ericsson et al. (2006); and *Development of Professional Expertise*, edited by Ericsson (2009).

The emphasis of this approach has shifted somewhat from general mechanisms of perception, memory, and problem-solving to knowledge-based and domain-specific issues of expertise. Although the classical well-constrained domains of chess and physics are still at the core of experimental research, more recent studies also include laboratory simulations and even observational field studies of real tasks of professional practice, chiefly in music, sports, medicine, law, and computer programming.

In practice, much of this work equates expertise with excellence. In their introductory chapter, Ericsson and Smith (1991) define the "original expertise approach" as seeking to "understand and account for what distinguishes outstanding individuals in a domain from less outstanding individuals in that domain" (p. 3). They point out that the approach focuses on those cases where the outstanding behavior can be attributed to "relatively stable characteristics of the corresponding individuals" (p. 3). The traditional study of expertise is basically the identification of superior and stable individual performances that are reproducible under standardized laboratory conditions. Given these requirements, it is no surprise that the most frequently studied form of expert performance is memory for meaningful stimuli from a well-constrained task domain (Ericsson & Smith, 1991, p. 23). Ericsson and Smith summarize the empirical findings of the human information-processing approach to expertise as follows:

> The superior performance consists of faster response times for the tasks in the domain, where we include the superior speed of expert typists, pianists, and Morse code operators. In addition, chess experts exhibit superior ability to plan ahead while selecting a move. … In a wide range of task domains experts have been found to exhibit superior memory performance. (p. 38)

In the overview of their volume, Glaser and Chi (1988, p. xvii–xx) summarized their view of the central findings of this approach in the form of seven points: (1) experts excel mainly in their own domain; (2) experts

perceive large meaningful patterns in their domain; (3) experts are fast: they are faster than novices at performing the skills of their domain, and they quickly solve problems with little error; (4) experts have superior short-term and long-term memory; (5) experts see and represent a problem in their domain at a deeper (more principled) level than novices do; novices tend to represent a problem at a superficial level; (6) experts spend a great deal of time analyzing a problem qualitatively; (7) experts have strong self-monitoring skills.

In contrast to the algorithmic approach, the enculturational approach to expertise sees thinking and knowledge as embedded in social situations, practices, and cultures. Knowledge and thought cannot be divorced from their corresponding skills and actions. As Collins (1987, p. 331) points out, "An apprenticeship, or at least a period of interpersonal interaction, is thought to be the necessary prelude to the transfer of skill-related knowledge." The mastery exhibited by an expert is above all tacit and intuitive. It is based on years of practical experience, not on the teaching of verbalized concepts and explicit algorithms. A strong formulation of this approach was put forward by Hubert and Stuart Dreyfus (1986) in their book *Mind over Machine*. An early collection of research within this approach may be found in the volume *Knowledge, Skill and Artificial Intelligence*, edited by Göranzon and Josefson (1988). Proponents of this approach seek philosophical support in the works of Polanyi and in late Wittgenstein (e.g., Nyíri & Smith, 1988).

The two approaches have often been presented as mutually exclusive rivals. There is, in fact, a very conspicuous aspect in which they seem to represent opposing views, namely the explicitness or verbalizability of expert thinking and knowledge. For Dreyfus and Dreyfus (1986, p. 30), "an expert's skill has become so much a part of him that he need be no more aware of it than he is of his own body." For Glaser and Chi (1988, p. xx), "Experts seem to be more aware than novices of when they make errors, why they fail to comprehend, and when they need to check their solutions." Dreyfus and Dreyfus see expert thinking as typically a non-symbolic process, whereas Glaser and others seem to take some sort of symbolization for granted.

However, this difference is less absolute than it first seems. Robert Hamm (1988) points out that the degree of explicitness and verbalization, as well as the use of analytical or intuitive modes of thinking, is dependent on the task at hand. Tasks of solitary problem-solving in a familiar domain are often accomplished without externally noticeable symbolic means. Tasks requiring negotiation and agreement among

members of a team can hardly be accomplished without some sort of explicit symbolic means.

Whatever importance the differences between the two approaches may have, their fundamental similarities are striking. These similarities have been largely overlooked in the literature, probably because they are mainly taken as self-evident assumptions by proponents of both approaches. They may be expressed in the form of three central propositions. I will formulate these three ideas polemically. The first part of each proposition is a positive statement; the latter part expresses a negative implication of the first part.

1. Locus of expertise: Expertise consists of superior and stable individual mastery of discrete tasks and skills. The understanding of expertise does not require that a more encompassing collective activity be taken as a unit of analysis.
2. Composition of expertise: Within a given domain of knowledge and practice, expertise is universal and homogeneous, and each domain is relatively self-sufficient. The aim is to identify "*the* expert" in a given field. There is no need to cross boundaries and build hybrids between substantively different types of expertise within the given domain, and collaboration across domains is not a core feature of – but an extension of – expertise.
3. Nature of learning involved in expertise: Expertise is acquired through internalization of experience, gained gradually by massive amounts of practice in the stable skills exhibited by the established masters of the given specialty (the famous novice–expert continuum). Expertise does not include questioning or reconceptualizing the skills and knowledge of established masters, nor the generation of culturally deviant and novel models of knowing and practice.

These three are core ideas of an individualist view, which depicts the mind as a solitary, self-sufficient mechanism (see Marková, 1982). Individualism goes hand in hand with an assumption of a stable status quo, a reluctance to focus on and conceptualize the creation of new culture as an ongoing collaborative achievement.

Serious problems in mainstream models of expertise began to surface in the 1980s and 1990s. A number of studies on expert decision-making found a pervasive tendency toward overconfidence and compartmentalization in the judgments of experts in various domains. Massive amounts of experience in no way guarantee an improved ability to deal with

uncertainty and probabilistic reasoning tasks (Brehmer, 1980). Experts often "appear to be mainly interested in how consistent the evidence is with the hypothesis they are testing and fail to consider its consistency with [an] alternative hypothesis" (Ayton, 1992, p. 95). Sternberg and Frensch (1992, p. 197) pointed out that "it is exceedingly difficult to break up and reorganize an automatized local processing system to which one in all likelihood no longer even has conscious access." In a similar vein, Argyris (1992) coined "skilled incompetence" as the dilemma of professionals. In an insightful early paper, Shchedrovitskii and Kotel'nikov (1988, p. 58) summed up the problem that was emerging:

> Today, in operating the technical systems we have created, and in the process of our ever-expanding appropriation of the world around us, we continually encounter assignments and tasks whose solution is beyond the capacities of any one person and requires the participation of a large team that includes representatives of different professions, different scientific disciplines, and different subjects. However, the coordinated organization of all these people into one working system has, as a rule, proved impossible: a person's thinking, organized by profession and subject, poses obstacles that are difficult to overcome, and a high level of professionalism interferes with, more than helps to achieve, joint team effort.

These critiques have continued and expanded. Hatchuel and Weil (1995) demonstrated the limits of traditional notions of expertise in conditions of continuous change in organizations. Faulkner, Fleck, and Williams (1998, p. 22) pointed out that "in order to make sensible and fair decisions, politicians, managers and (most of all) citizens need to draw not only on the expertise and tools of scientists and technologists, but also on crucial social and economic knowledge which technical people generally lack and/or undervalue." Martin (1996) presented a set of cases in which people confronted established experts. Selinger and Crease (2006) published a collection of philosophical discussions of the limits of expertise. Much of the dissatisfaction with dominant views of expertise was summed up by Freedman (2010) in his book *Wrong: Why Experts Keep Failing Us – and How to Know When Not to Trust Them*.

Finally, there is the ever-popular genre, spanning the years from Illich (1973) to Susskind and Susskind (2015), of critiquing dominant forms of expertise as restrictive monopolies of self-serving professions. Suggested solutions range from Illich's preindustrial "tools for conviviality" to the supposedly liberating impact of postindustrial digital technologies, promising to make expert knowledge accessible to everybody. Predictably, these critiques have generated a countercritique that argues

that the denigration of expertise is a threat that leads to populism and irrationalism (Nichols, 2017).

### Going Beyond Individualism and Stability in the Study of Expertise

In the early 1990s, Jean Lave and Etienne Wenger (1991) opened what was to become a multifaceted stream of discussions and studies on situated learning. They suggested that the proper unit of analysis of skilled human activity is a *community of practice* rather than an isolated individual. Skill, knowledge, and competence reside in local working communities, not in transportable packages or in the heads of individual subjects. They also suggested that the foundational mechanism of becoming competent in a domain is *legitimate peripheral participation* in a relevant community of practice rather than transmission of knowledge in school-like forms. Legitimate peripheral participation may best be observed in various settings of apprenticeship.

Two years later, Carl Bereiter and Marlene Scardamalia (1993) published a book titled *Surpassing Ourselves*. They also criticized strictly individualist notions of expertise and suggested that teamwork should be taken seriously as a variety of expertise. More important, they suggested that expertise should be reconceptualized as a process of going beyond the normal course of learning, as *progressive problem-solving*. According to these authors, experts "tackle problems that increase their expertise, whereas nonexperts tend to tackle problems for which they do not have to extend themselves" (Bereiter & Scardamalia, 1993, p. 78). Instead of trying to reduce novel problems to simple components that can be handled with familiar routine procedures, experts construct new concepts and methods for unfamiliar cases.

In 1995, Edwin Hutchins published a book titled *Cognition in the Wild*. He maintained that cognition in real-world settings is typically not a solitary achievement of an individual but a *distributed* achievement of a *functional system* consisting of human practitioners, their artifacts, and their representations. Cognitive performance such as expert problem-solving is best analyzed as the propagation of representational states across humans and artifacts in a functional system, for example, in a unit responsible for the navigation of a large ship or in the cockpit of a passenger jet. The acquisition of expertise takes place as members of such distributed functional systems gradually acquire a broader and more flexible mastery of the task domain for which the system is responsible, and as the system itself adapts to changing circumstances.

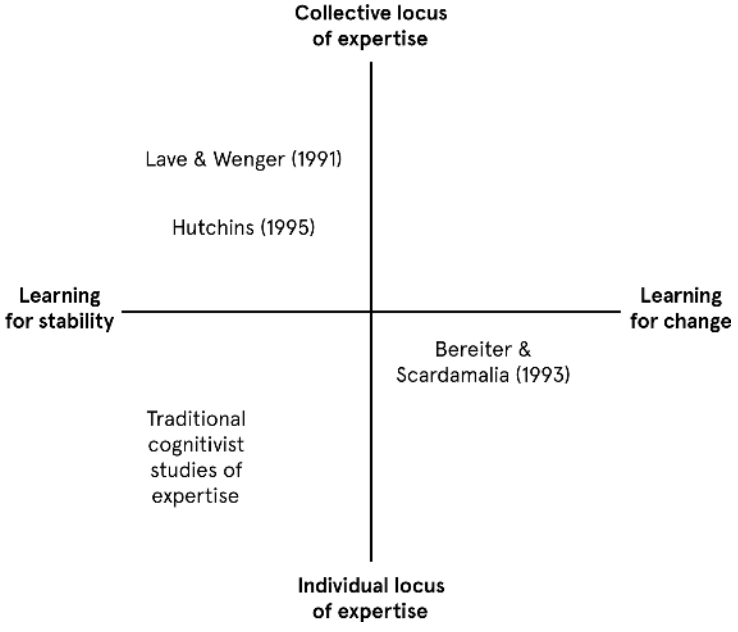


FIGURE 1.1. Early landmarks in practice-based studies of expertise.

The three books are important landmarks of a wave of research and theorizing that opened up a new practice-based perspective on expert work and cognition. This wave was continued and enriched by Engeström and Middleton (1996), Keller and Keller (1996), Grint and Woolgar (1997), Wenger (1998), Heath and Luff (2000), John-Steiner (2000), Luff, Hindmarsh, and Heath (2000), Engeström, Lompscher, and Rückriem (2005), and others. We might talk about a turn to collective, culturally situated practices of expertise.

The contributions of Lave and Wenger, Bereiter and Scardamalia, and Hutchins may be characterized with the help of a two-dimensional conceptual space depicted in Figure 1.1. The vertical dimension represents the locus of expertise, ascending from the traditional cognitivist notion of the sphere of an isolated individual to the sphere of a team or functional system, to a community of practice – and potentially all the way up to a field of multiple interacting communities dealing with partially shared objects and tasks.

Along this vertical dimension, Bereiter and Scardamalia (1993) stayed closest to the traditional emphasis on the individual expert. Hutchins (1995) focused on relatively well-bounded functional systems or teams,